This document is scheduled to be published in the
Federal Register on 09/11/2024 and available online at
https://federalregister.gov/d/2024-20529, and on https://govinfo.gov

0-JT-P

DEPARTMENT OF COMMERCE

Bureau of Industry and Security

15 CFR Part 702

[Docket No. 240905-0231]

RIN 0694-AJ55

Establishment of Reporting Requirements for the Development of Advanced Artificial Intelligence Models and Computing Clusters

AGENCY: Bureau of Industry and Security, Department of Commerce.

ACTION: Proposed rule; request for comment

SUMMARY: This proposed rule would amend the Bureau of Industry and Security's (BIS) Industrial Base Surveys – Data Collections regulations by establishing reporting requirements for the development of advanced artificial intelligence (AI) models and computing clusters under the Executive order of October 30, 2023, "Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence."

DATES: Comments on this proposed rule must be received by BIS by no later than [INSERT 30 DAYS FROM DATE OF PUBLICATION IN THE FEDERAL REGISTER].

ADDRESSES: Comments on this proposed rule may be submitted to the Federal rulemaking portal (www.regulations.gov). The regulations.gov ID for this proposed rule is: BIS–2024–0047. Please refer to RIN 0694–AJ55 in all comments.

Anyone submitting business confidential information should clearly identify any business confidential portion of a comment at the time of submission, file a statement justifying nondisclosure and referring to the specific legal authority claimed, and provide a non-confidential version of the submission.

For comments submitted electronically containing business confidential information, the file name of the business confidential version should begin with the characters “BC.” Any page containing business confidential information must be clearly marked “BUSINESS CONFIDENTIAL” on the top of that page. The corresponding non-confidential version of those comments must be clearly marked “PUBLIC.” The file name of the non-confidential version should begin with the character “P.” Any submissions with file names that do not begin with either a “BC” or a “P” will be assumed to be public and will be made publicly available through https://www.regulations.gov. Commenters submitting business confidential information are encouraged to scan a hard copy of the non-confidential version to create an image of the file, rather than submitting a digital copy with redactions applied, to avoid inadvertent redaction errors which could enable the public to read business confidential information.

FOR FURTHER INFORMATION CONTACT: Sean Delehanty, Office of Strategic Industries and Economic Security Bureau of Industry and Security, Department of Commerce. Phone: 202-316-5765; Email: Sean.Delehanty@bis.doc.gov.

SUPPLEMENTARY INFORMATION:

Background

Section 4.2(a)(i) of Executive Order 14110 of October 30, 2023, “Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence” (E.O. 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. Additionally, section 4.2(a)(ii) of E.O. 14110 directs the Secretary 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 E.O. 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 reporting requirements proposed in this regulation are intended to apply to dual-use foundation models that meet technical conditions issued by the Department. The Department expects to update the technical conditions, based on technological advancements, as necessary and appropriate, as directed by section 4.2(b) of E.O. 14110.

E.O. 14110 directs the Department of Commerce (Department) to collect information about dual-use foundation models in accordance with the Defense Production Act (DPA) (50 U.S.C. 4501 *et seq*.). Under the DPA, the President is authorized to take actions that ensure the U.S. industrial base is prepared to supply products and services to support the national defense. In this context, the term “national defense” means “programs for military and energy production or construction, military or critical infrastructure assistance to any foreign nation, homeland security, stockpiling, space, and any directly related activity” (50 U.S.C. 4552(14)). Additionally, the DPA makes clear that the international competitiveness of the U.S. industrial base directly affects its ability to support the national defense (see 50 U.S.C. 4502(a)(7)).

Among other authorities, the DPA authorizes the President “by regulation, subpoena, or otherwise, to obtain such information from, require such reports and the keeping of such records by, make such inspection of the books, records, and other writings, premises or property of, and take the sworn testimony of, and administer oaths and affirmations to, any person as may be necessary or appropriate, in his discretion, to the enforcement or the administration of” the DPA (50 U.S.C. 4555(a)). The DPA further specifies that this grant of authority “includes the authority to obtain information in order to perform industry studies assessing the capabilities of the United States industrial base to support the national defense" (50 U.S.C. 4555(a)).

To carry out its obligations under section 4.2(a) of E.O. 14110, BIS is exercising its DPA authority, which was delegated to the Department by the President in Executive Order 13603, and subsequently re-delegated within the Department to BIS, to collect information from U.S. companies that are developing, have plans to develop, or have the computing hardware necessary to develop dual-use foundation models. AI models are quickly becoming integral to numerous U.S. industries that are essential to the national defense. For example, manufacturers of military equipment (*e.g.*, aircrafts, tanks, and missile launchers) use AI models to enhance the maneuverability, accuracy, and efficiency of equipment.[1] Similarly, manufacturers of signal intelligence devices (*e.g.*, satellites, cameras, and radar) use AI models to improve how those devices capture signals and eliminate noise.[2] As a final example, developers of cybersecurity software, which can be applied to protect a wide range of systems and infrastructure that are critical to the national defense, use AI models to increase the speed at which that software detects and responds to cyberattacks.[3]

Dual-use foundation models could increase the capabilities of these products and services to an even greater extent. Specifically, integrating dual-use foundation models into products like military equipment, signal intelligence devices, and cybersecurity software could enable those

[1] Shield AI, "Shield AI Conducts AI-Piloted Flights on Sixth Aircraft, the Kratos MQM-178 Firejet" (Mar. 29, 2024), https://shield.ai/shield-ai-conducts-ai-piloted-flights-on-sixth-aircraft-the-kratos-mqm-178-firejet/; RTX, "Raytheon Technologies Unveils Next-Generation Electro-Optical Intelligent-Sensing Capability (Apr. 24, 2023), https://www.rtx.com/news/news-center/2023/04/24/raytheon-technologies-unveils-next-generation-electro-optical-intelligent-sensing.

[2] National Instruments, "Artificial Intelligence in Software-Defined SIGINT Systems" (Feb. 6, 2024), https://www.ni.com/en/solutions/aerospace-defense/radar-electronic-warfare-sigint/artificial-intelligence-in-software-defined-sigint-systems.html; Northrop Grumman, "Artificial Intelligence Helps Protect Troops in Denied GPS Environments", https://www.northropgrumman.com/what-we-do/artificial-intelligence-helps-protect-troops-in-denied-gps-environments.

[3] NVIDIA, "NVIDIA and Booz Allen Hamilton Expand Partnership to Bring AI-Enabled Cybersecurity to Public and Private Sectors" (Sept. 20, 2022), https://nvidianews.nvidia.com/news/nvidia-and-booz-allen-hamilton-expand-partnership-to-bring-ai-enabled-cybersecurity-to-public-and-private-sectors; IBM, "AI and Automation for Cybersecurity" (June 2022), https://www.ibm.com/thought-leadership/institute-business-value/en-us/report/ai-cybersecurity.

products to operate more effectively across a wider range of environments, to respond more effectively to unexpected signals, and to combat additional types of cyberattacks.

Given those potential capabilities, it is essential to the national defense that the defense industrial base is able to integrate dual-use foundation models. Indeed, because industries and governments across the world are actively working to integrate dual-use foundation models into their defense capabilities, the U.S. defense industrial base will need to integrate dual-use foundation models to remain internationally competitive.

Accordingly, the U.S. Government must be ready to take actions that ensure dual-use foundation models produced by U.S. companies are available to the defense industrial base. To do so, the U.S. Government needs information about how many U.S. companies are developing, have plans to develop, or have the computing hardware necessary to develop dual-use foundation models, as well as information about the characteristics of dual-use foundation models under development. Such information will allow the U.S. Government to determine whether action is necessary to stimulate development of dual-use foundation models or to support the development of specific types of models.

The integration of AI models into the defense industrial base also requires the U.S. Government to take actions as needed to ensure that dual-use foundation models operate in a safe and reliable manner. Products integrating these models may operate in unpredictable or unreliable ways, potentially resulting in dangerous accidents, and a lack of reliability will make it difficult for the U.S. Government to use those products in contexts where the margin for error is small, including defense-related activities where accidents could result in injury or even loss of life. Thus, the U.S. Government needs information about how companies developing dual-use foundation models are training those models to respond to different kinds of inputs and information about how those companies have tested the safety and reliability of their models. Such information will allow the U.S. Government to determine the extent to which certain dual-use foundation models can be used by the defense industrial base and whether action is needed to ensure that the defense industrial base produces the safest and most reliable products and services in the world.

For similar reasons, the U.S. Government must minimize the vulnerability of dual-use foundation models to cyberattacks. Dual-use foundation models can potentially be disabled or manipulated by hostile actors, and it will be difficult for the U.S. Government to rely on a particular model unless it can determine that the model is robust against such attacks. Accordingly, the U.S. Government needs information about the cybersecurity measures that companies developing dual-use foundation models use to protect those models, as well as information about those companies' cybersecurity resources and practices. Under 15 CFR 702.3 all information submitted to the Department under this rule will be treated as confidential and afforded all the protections of section 705(d) of the DPA. Such information will allow the U.S. Government to determine which models are secure enough to be integrated into products or services that are essential to the national defense and to assess whether action is needed to ensure that the defense industrial base is producing the most secure products and services in the world.

Finally, the U.S. Government must prepare the defense industrial base for the possibility that foreign adversaries or non-state actors will use dual-use foundation models for activities that threaten the national defense, including to develop weapons and other dangerous technologies. Accordingly, the U.S. Government requires information about the safety and reliability of AI models, including any potentially dangerous capabilities that developers of dual-use foundation models have identified with respect to those models. This includes the results of tests related to reliability as well as the results of any red-team testing that the company has conducted relating to lowering the barrier to entry for the development, acquisition, and use of biological, weapons by non-state actors; the discovery of software vulnerabilities and development of associated exploits; the use of software or tools to influence real or virtual events; and the possibility for self-replication or propagation. Such information will enable the U.S. Government to determine whether investments in the defense industrial base are needed to ensure the United States has access to safe and reliable AI systems, as well as to counteract the dangerous capabilities identified or to ensure that adequate safeguards are in place to prevent the theft or misuse of dual-use foundation models by foreign adversaries or non-state actors.

In short, dual-use foundation models will likely drive significant advances in numerous industries on which the national defense depends. These advances require BIS to conduct an ongoing assessment of the AI industry to ensure that the U.S. Government has the most accurate, up-to-date information when making policy decisions about the international competitiveness of the industrial base and its ability to support the national defense.

Section 4.2(a)(i) of E.O. 14110 mandates that the Secretary shall require companies developing dual-use foundation AI models to provide information, reports, or records regarding the following:

1. any ongoing or planned activities related to training, developing, or producing dual-use foundation models, including the physical and cybersecurity protections taken to assure the integrity of that training process against sophisticated threats;

2. the ownership and possession of the model weights of any dual-use foundation models, and the physical and cybersecurity measures taken to protect those model weights; and

3. the results of any developed dual-use foundation model's performance in relevant AI red-team testing, including a description of any associated measures the company has taken to meet safety objectives, such as mitigations to improve performance on these red-team tests and strengthen overall model security.

4. Other information pertaining to the safety and reliability of dual-use foundation models, or activities or risks that present concerns to U.S. national security.

Section 4.2(a)(ii) of EO 14110 also mandates that companies, individuals, or other organizations or entities that acquire, develop, or possess a potential large-scale computing cluster must 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. To the extent that these entities are companies developing dual-use foundation models, they are also subject to obligations 1-3, above.

Discussion of the Proposed Rule

This proposed rule outlines a potential notification and reporting process for companies developing or intending to develop dual-use foundation AI models and for companies, individuals or other organizations or entities that acquire, develop, or possess computing clusters that meet technical conditions issued by the Department. Such entities would be required to report the required information to the BIS on a quarterly basis for activities that occurred during that quarter or that are planned to occur in the six months following the quarter.

BIS collected information responsive to the requirements of section 4.2(a) of E.O. 14110 via a mandatory survey of companies identified as developing or planning to develop potential dual-use foundation models. That survey was issued on January 26, 2024. Under this proposed rule, companies that completed the survey and any other companies that have developed or are in the process of developing dual-use foundation models or large-scale computing clusters would be required to submit information about these activities on a quarterly basis.

For companies that have already submitted complete information via the survey, the reporting requirements will not require that the company report activity already reported to BIS in the survey but would require the reporting of any additions, updates, or changes to the information since the survey. Any company that has filed at least one report would be required to continue to file reports on a quarterly basis for as long as it continues to meet the reporting requirements or, if it no longer meet the requirements, until it has filed seven quarterly reports affirming that it has no additions, updates, or changes to the information in the last report. The reporting system will allow for companies that have no additions, updates, or changes since the last report to make a simple notification to that effect.

Request for Comments

BIS welcomes public comment on all aspects of this proposed regulation. While much of the information that entities must report is dictated by section 4.2(a) of E.O. 14110, BIS is particularly interested in public comments on the following:

1. Quarterly Notification Schedule: BIS has proposed that all covered U.S. persons with models or clusters exceeding the technical thresholds for reporting should notify BIS on a quarterly basis. Covered U.S. persons would be required to make quarterly notifications of 'applicable activities' that meet the criteria under § 702.7(a)(1)(i) or (ii) planned to occur in the next six months related to dual-use foundation models and/or computing clusters, as well as quarterly notifications required for any 'applicable activities' (*i.e.*, an "applicable activity" that meets the criteria under § 702.7(a)(1)(i) or (ii)) and § 702.7(a)(2)(v) (*Affirmation of no applicable activities*), as applicable. 'Applicable activities' are defined to include developing, or having the intent to develop within the next six months, an AI model or computing cluster above certain technical thresholds specific in this proposed rule. If a covered U.S. person has any 'applicable activities' to report, then they will notify BIS, and BIS will follow up with more detailed questions, to which the Covered U.S. person must respond within 30 calendar days. If Covered U.S. persons have no 'applicable activities' to report, they would only be required to affirm that fact to BIS each quarter. BIS has proposed a quarterly notification schedule to provide the U.S. Government with timely information on the safety and security of large AI models and computing clusters, while offering a regular notification schedule to facilitate respondent planning and ease respondent burden. BIS welcomes comments on the frequency of the proposed notification schedule, as well as alternatives for achieving timely reporting of the required information.

2. Collection and Storage: BIS recognizes that the information collected through these reporting requirements is extremely sensitive. In the interest of gathering information on prioritizing the safety of respondents' data, BIS welcomes comments related to how this data should be collected and stored.
3. Collection Thresholds: BIS has included the technical conditions specified in E.O. 14110 for models and computing clusters that would trigger the proposed reporting requirements. As directed by section 4.2(b) of E.O. 14110, BIS will update these technical conditions as appropriate. In addition to the technical parameters in E.O. 14110, BIS is also seeking comments on the following proposed updated collection parameters. BIS welcomes comments on the following sets of technical parameters.
   - A dual-use foundation model training run triggers reporting requirements if it utilizes more than 10^26 computational operations (*e.g.*, integer or floating-point operations). Models trained on primarily biological sequence data, but at the lower threshold of 10^23 computational operations, as specified by section 4.2(b) of E.O. 14110, will be addressed in a separate survey.
   - Large-scale computing clusters are defined as clusters having a set of machines transitively connected by networking of over 300 Gbit/s and having a theoretical maximum performance greater than 10^20 computational operations (*e.g.*, integer or floating-point operations) per second (OP/s) for AI training, without sparsity.

Rulemaking Requirements

1. This proposed rule has been determined to be a significant regulatory action for purposes of E.O. 12866.

2. Notwithstanding any other provision of law, no person is required to respond to, nor shall any person be subject to a penalty for failure to comply with, a collection of information subject to

the requirements of the Paperwork Reduction Act of 1995 (44 U.S.C. 3501 *et seq.*) (PRA), unless that collection of information displays a currently valid Office of Management and Budget (OMB) Control Number. This proposed rule involves a currently approved information collection *National Security and Critical Technology Assessments of the US Industrial Base* (OMB Control Number 0694-0119). The authority for this collection is section 705 of the Defense Production Act of 1950, as amended and related Executive Orders 12656 and 13603. Under this information collection, BIS conducts surveys and assessments of critical U.S. industrial sectors and technologies. Undertaken at the request of various policy, research and development, and program and planning organizations within the Department of Defense and the Armed Services, Department of Homeland Security (DHS), National Aeronautics and Space Administration (NASA), and other agencies, BIS research, data collection and analysis provide needed information to benchmark industry performance and raise awareness of diminishing manufacturing capabilities.

Most surveys include questions necessary to obtain data on employment, supply chain, financial performance, production, technology and service capabilities, research and development (R&D), investment, competitive outlook, export controls and other relevant information. Some surveys include a few non-standard questions, depending on the industry and the needs of the partner agency. The number of surveys required per assessment varies with the size of the sector and the scope of the project.

Information gathered from these surveys is deemed business confidential and will be treated in accordance with section 705 of the Defense Production Act of 1950 which prohibits the publication or disclosure of such information unless the President determines that its withholding is contrary to the national defense. To review previous surveys cleared under this generic collection—including all background materials—please visit at *https://www.reginfo.gov/public/do/PRAMain* and use the search function to enter either the title of the collection or the OMB Control Number.

When this proposed rule is finalized, BIS intends to use this existing information collection for the collection/reporting requirement required by E.O. 14110. BIS estimates the specific survey required by this proposed rule will have an estimated burden of 5,000 hours per year aggregated across all new respondents. BIS believes this increase in respondent burden does not require a change to the burden or cost estimates for the overall umbrella clearance. Please see the request for comment section of the proposed rule for more information the potential information collection elements BIS is considering for the final rule and subsequent surveys.

3. These proposed changes do not contain policies with federalism implications as that term is defined in EO 13132.

4. The Regulatory Flexibility Act (RFA), as amended by the Small Business Regulatory Enforcement Fairness Act of 1996 (SBREFA) (5 U.S.C. 601 *et seq.*) generally requires an agency to prepare a regulatory flexibility analysis of any rule subject to the notice and comment rulemaking requirements under the Administrative Procedure Act (5 U.S.C. 553) or any other statute. Under section 605(b) of the RFA, however, if the head of an agency certifies that a rule will not have a significant impact on a substantial number of small entities, the statute does not require the agency to prepare a regulatory flexibility analysis. Pursuant to section 605(b), the Chief Counsel for Regulation, Department of Commerce, certified to the Chief Counsel for Regulation, Small Business Administration that this proposed rule will not have a significant impact on a substantial number of small entities for the reasons explained below. No other law

requires such an analysis. Consequently, no regulatory flexibility analysis is required, and none has been prepared.

Number of Small Entities

Small entities include small businesses, small organizations, and small governmental jurisdictions. For purposes of assessing the impacts of this proposed rule on small entities, a small business, as described in the Small Business Administration's Table of Small Business Size Standards Matched to North American Industry Classification System (NAICS) Codes (effective March 17, 2023), has a maximum annual revenue of $47 million and a maximum of 1,500 employees (for some business categories, these numbers are lower). A small governmental jurisdiction is a government of a city, town, school district or special district with a population of less than 50,000. A small organization is any not-for-profit enterprise which is independently owned and operated and is not dominant in its field. The most apt code to apply here is NAICS 518 - Computing Infrastructure Providers, Data Processing, Web Hosting, and Related Services.

The reporting requirements in this proposed rule are expected to apply to only a small number of entities – only those companies developing or intending to develop a dual-use foundation model and those companies, individuals, or other organizations or entities that acquire, develop, or possess potential large-scale computing clusters. For the purposes of this rulemaking, the term "covered U.S. persons" includes all U.S. persons subject to the reporting requirements of E.O. 14110, section 4.2(a), and is defined as any individual U.S. citizen, any lawful permanent resident of the United States as defined by the Immigration and Nationality Act, any entity—including organizations, companies, and corporations—organized under the laws of the United States or any jurisdiction within the United States (including foreign branches), or any person (individual) located in the United States. At present, BIS assesses that there are between zero and 15 companies exceed the reporting thresholds for models and computing clusters at the time of publication. All of these entities are well-resourced technology companies. Exceeding the technical thresholds for models and computing clusters requires access to vast computing power, which is not typically available to small entities. The minimum computational threshold that would trigger a reporting requirement established in E.O. 14110 currently exceeds all or virtually all models in use.[4]

As AI technology development and implementation are expected to advance over the next few years, the number of covered U.S. persons involved in it will also increase. However, as directed by E.O. 14110, the Secretary will update the technical conditions that trigger the reporting requirements over time, which may limit the number of additional impacted entities over time.

Impact

For the reasons discussed above, BIS believes that this proposed rule, which would impose reporting requirements on large technology companies, would have no significant impact on small entities.

Conclusion

BIS believes that the overall impact of this proposed rule on small entities would not be significant, as it would only apply to entities with large monetary and computational resources, which BIS believes are not small entities. For the reasons set forth above, the Chief Counsel for Regulations at the Department of Commerce has certified that this action would not have a significant impact on a substantial number of small entities.

[4] Rahman, Owen, and You. “Tracking Large-Scale AI Models” (April 5, 2024), https://epochai.org/blog/tracking-large-scale-ai-models.

In accordance with 5 U.S.C. 553(b)(4), a summary of this proposed rule may be found at www.regulations.gov. The regulations.gov ID for this proposed rule is: BIS–2024–0047.

List of Subjects in 15 CFR Part 702

Business and industry, Confidential business information, Employment, National defense, Penalties, Research, Science and technology.

Accordingly, 15 CFR part 702 is proposed to be amended as follows:

PART 702 – INDUSTRIAL BASE SURVEYS – DATA COLLECTIONS

1. The authority citation for 15 CFR part 702 is revised to read as follows:

   Authority: 50 U.S.C. 4501 *et seq.*; E.O. 13603, 77 FR 16651, 3 CFR, 2012 Comp., p. 225; E.O. 14110, 88 FR 75191, 3 CFR, 2023 Comp., p. 657.

2. Section 702.7 is added to read as follows:

§ 702.7 Special requirements for on-going reporting regarding the development of advanced artificial intelligence models and computing clusters.

(a) *Reporting requirements.* (1) Covered U.S. persons are required to submit a notification to the Department by emailing <ai_reporting@bis.doc.gov> on a quarterly basis as defined in paragraph (a)(2) of this section if the covered U.S. person engages in, or plans, within six months, to engage in ‘applicable activities,’ defined as follows:

(i) Conducting any AI model training run using more than 10^26 computational operations (*e.g.*, integer or floating-point operations); or

(ii) Acquiring, developing, or coming into possession of a computing cluster that has a set of machines transitively connected by data center networking of greater than 300 Gbit/s and having a theoretical maximum greater than 10^20 computational operations (*e.g.*, integer or floating-point operations) per second (OP/s) for AI training, without sparsity.

Note 1 to paragraph (a)(1): Consistent with industry conventions, one multiply-accumulate computation, D=AxB+C, should be counted as two operations.

(2) *Timing of notifications and response to BIS questions*—(i) *Notification of applicable activities*. Covered U.S. persons subject to the reporting requirements in paragraph (a)(1) of this section must notify BIS of 'applicable activities' via email each quarter, identifying any 'applicable activities' planned in the six months following notification. Quarterly notification dates are as follows: Q1- April 15; Q2- July 15; Q3- October 15; Q4- January 15. For example, in a notification due on April 15, a covered U.S. person should include all activities planned until October 15 of the same year.

(ii) *Response to BIS questions.* Following a notification of 'applicable activities' by a covered U.S. person, the covered U.S. person will receive questions from BIS. The covered U.S. person must respond to all questions within 30 calendar days of receiving the request.

(iii) *Corrections.* If any notification of 'applicable activities' or response to BIS questions filed under this section is incomplete when filed, BIS will notify the covered U.S. person and require a revised resubmission within 14 calendar days after BIS provides notice of incompletion. BIS will continue to require revisions within 14 calendar days of notification if a resubmission remains incomplete.

(iv) *Clarification questions.* If, after receipt of responses described in paragraph (a)(2)(ii) of this section, BIS has additional questions to clarify those responses, the covered U.S. person will provide additional responses to such additional questions within seven (7) calendar days. If the covered U.S. person needs additional time to provide an additional response, it can request an extension from BIS.

(v) *Affirmation of no applicable activities.* For each of the seven quarters following the quarter covered by a notification of 'applicable activities,' if the covered U.S. person has no 'applicable activities' (*i.e.*, an "applicable activity" that meets the criteria under paragraph (a)(1)(i) or (ii) of this section) to report, they must submit an affirmation of no applicable activities by emailing <ai_reporting@bis.doc.gov> on the quarterly notification date. If the covered U.S. person submits an affirmation of no applicable activities for seven consecutive quarters, they need not provide BIS with any affirmation thereafter until they have 'applicable activities' to report.

(b) *Content, form, and manner of response to BIS questions.* (1) All information submitted under this section shall be filed with BIS in the form and manner that BIS will prescribe in instructions sent to the covered U.S. person after BIS has received a notification of 'applicable activities.'
(2) BIS will send questions to the covered U.S. person which must address, but may not be limited to, the following topics:
(i) Any ongoing or planned activities related to training, developing, or producing dual-use foundation models, including the physical and cybersecurity protections taken to assure the integrity of that training process against sophisticated threats;

(ii) The ownership and possession of the model weights of any dual-use foundation models, and the physical and cybersecurity measures taken to protect those model weights;

(iii) The results of any developed dual-use foundation model's performance in relevant AI red-team testing, including a description of any associated measures the company has taken to meet safety objectives, such as mitigations to improve performance on these red-team tests and strengthen overall model security; and

(iv) Other information pertaining to the safety and reliability of dual-use foundation models, or activities or risks that present concerns to U.S. national security.

(c) *Definitions.* For purposes of the reports required by paragraph (a) of this section, apply the following definitions.

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.

AI model means a component of an information system that implements AI technology and uses computational, statistical, or machine-learning techniques to produce outputs from a given set of inputs.

AI system means any data system, software, hardware, application, tool, or utility that operates in whole or in part using AI.

Artificial intelligence or *AI* has the meaning set forth in 15 U.S.C. 9401(3).

Company means a corporation, partnership, association, or any other organized group of persons, or legal successor or representative thereof. This definition is not limited to commercial or for-profit organizations. For example, the term "any other organized group of persons" may encompass academic institutions, research centers, or any group of persons who are organized in some manner. The term "corporation" is not limited to publicly traded corporations or corporations that exist for the purpose of making a profit.

Covered U.S. person means any individual U.S. citizen, lawful permanent resident of the United States as defined by the Immigration and Nationality Act, entity—including organizations, companies, and corporations—organized under the laws of the United States or any jurisdiction within the United States (including foreign branches), or any person (individual) located in the United States.

Dual-use foundation model means an AI model that is:

(i)(A) Trained on broad data;

(B) Generally uses self-supervision;

(C) Contains at least tens of billions of parameters;

(D) Is applicable across a wide range of contexts; and

(E) 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, such as by:

(*1*) Substantially lowering the barrier of entry for non-experts to design, synthesize, acquire, or use chemical, biological, radiological, or nuclear (CBRN) weapons;

(*2*) Enabling powerful offensive cyber operations through automated vulnerability discovery and exploitation against a wide range of potential targets of cyberattacks; or

(*3*) Permitting the evasion of human control or oversight through means of deception or obfuscation.

(ii) Models meet this definition even if they are provided to end users with technical safeguards that attempt to prevent users from taking advantage of the relevant unsafe capabilities.

Knowledge has the meaning set out in 15 CFR 772.1.

Large-scale computing cluster means a cluster of computing hardware that meets the technical thresholds provided by the Department in paragraph (a)(1) of this section.

Model weights means the numerical parameters used in the layers of a neural network.

Training or *training run* refers to any process by which an AI model learns from data using computing power. Training includes but is not limited to techniques employed during pretraining like unsupervised learning and employed during fine tuning like reinforcement learning from human feedback.

United States (U.S.) includes the 50 states, the District of Columbia, Puerto Rico, Guam, American Samoa, the U.S. Virgin Islands, and the Northern Mariana Islands.

Thea D. Rozman Kendler,

Assistant Secretary for Export Administration

[FR Doc. 2024-20529 Filed: 9/9/2024 8:45 am; Publication Date: 9/11/2024]