VedderPrice

President Biden Issues Far-Reaching Executive Order on Artificial Intelligence

By Daniel H. Shulman and Sudip K. Mitra

November 3, 2023

President Biden issued an Executive Order on October 30, 2023 designed to place the United States at the forefront of law and regulation of Artificial Intelligence (AI). The Executive Order on the "Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence" creates binding disclosure requirements for companies that are either developing certain large language AI models or acquiring or possess sufficient computing power to run such AI implementations (as described below). The Order also establishes, and directs several federal agencies to establish, industry benchmarks for ensuring robust, reliable, repeatable and standardized testing and evaluations of AI systems, create new standards for AI safety and security.

The Order contains a lot of detailed provisions and initiatives involving nearly every government agency and calling for wide-ranging studies and recommendations on nearly every facet of AI, significant provisions of which are described below.

Of particular note, however, the President invoked the Defense Production Act to impose certain requirements that will go into effect 90 days after the issuance of the Order. There are two significant requirements going into effect affecting companies that employ AI models and companies that employ or provide large computing capacity that can be used for AI.

First, the Department of Commerce will require companies employing certain large language models to provide, on an ongoing basis, information, records and reports on their training of the tools; the ownership and possession of the training models; and the results of testing for certain high risk AI implementations based on guidance provided by NIST. Companies are also required to provide description of any associated measures taken to meet safety objectives as well as mitigation procedures to improve performance and strengthen overall model security for the AI systems.

Second, the Order requires companies to report on their acquisition or possession of large clusters of computing power. Companies that acquire, develop, or possess a large-scale computing cluster must also report to the Commerce Department on certain factors, including (a) the existence and location of such clusters and (b) the amount of total computing power available in each such cluster, which could be potentially burdensome for some entities.

The Order's immediate reporting requirements apply to any large language model that it defines as a "dual-use foundation model," which is "an AI model that 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," with a list of examples of ways those risks could arise.

The requirements also, at least until the agencies come up with their own definitions, will apply to certain minimum levels of computing power, namely, "(i) any model that was trained using a quantity of computing power greater than 10^{26} integer or floating-point operations, or using primarily biological sequence data and using a quantity of computing power greater than 10^{23} integer or floating-point operations; and (ii) any computing cluster that has a set of machines physically co-located in a single datacenter, transitively connected by data center networking of over 100 Gbit/s, and having a theoretical maximum computing capacity of 10^{20} integer or floating-point operations per second for training AI."

Below are some of the other key aspects of the Order governing the standards and best practices for development of safe and secure AI systems and the agency responsibilities for the same.

The Order directs the National Institute of Standards and Technology (NIST), Department of Energy (DoE) and Department of Homeland Security (DHS) to establish and develop guidelines and best practices with the aim of promoting consensus industry standards for developing and deploying safe, secure and trustworthy AI systems within a period of 270 days from the date of the Order. Specifically, the Order creates agency-specific directives for generative AI and large language models and creates guidelines for the testing of AI in order to support security and trustworthiness and implement a plan for developing the Energy Department's AI model evaluation tools and testbeds, especially (and at a minimum) to evaluate AI systems' capabilities to proliferate weapons, infrastructure and energy-security threats. The Order mandates the NIST develop companion resources to incorporate secure development practices for generative AI and for dual-use foundation models (as detailed above) via conducting AI red-teaming tests and creates guidance for evaluating and auditing AI capabilities.

The Order prioritizes focusing on capabilities through which AI could cause harm, such as in the areas of cybersecurity and biosecurity. For instance, the Order mandates agencies to propose regulations requiring any U.S. Infrastructure as a Service (IaaS) providers to submit a report to the Secretary of Commerce when a foreign person transacts with that) provider to train a large AI model with potential capabilities that could be used in malicious cyber-enabled activity. More importantly, it also prohibits any foreign reseller of their U.S. IaaS product from providing those products unless such foreign reseller submits to the U.S. IaaS provider a report in compliance with the agency regulations. Companies in the United States associated with such foreign resellers will need to comply with such requirements under the Order.

The Order requires that the heads of each agency with relevant regulatory authority over critical infrastructure and the heads of relevant sector risk management agencies (SRMAs) must provide to DHS an assessment of potential risks related to the use of AI in critical-infrastructure sectors. For instance, within 150 days of issuance of the Order, the Secretary of Treasury is required to issue a public report on best practices for financial institutions to manage AI-specific cybersecurity risks.

The Order mandates the creation, within DHS, of an Artificial Intelligence Safety and Security Board as an advisory committee to DHS to advise on use of AI in critical infrastructure. Government agencies are tasked with evaluating and reporting on threats caused by AI, especially around energy, weapons (especially chemical, biological and nuclear), defense and government secrets. The State Department, Office of Science and Technology (OST), Department of Health and Human Services (HHS), DHS and Director of National Intelligence and Energy are required to examine the use of AI to identify biological (including DNA) sequences and methods of synthesis that would pose a risk to national security. The Commerce Department is also required to study and evaluate the pros and cons of having model weights (e.g., algorithms) made public, or removing safeguards from publication of those models, and to recommend regulatory mechanisms around widely available model weights.

Other aspects of the Order touch on AI's influence in the realm of intellectual property. First, within 270 days of issuance of the order, both the U.S. Patent and Trademark Office (USPTO) and U.S. Copyright Office are required to issue reports on AI, including updated guidance on patent and copyright eligibility for AI products and systems. The Order also requires the Commerce Department to develop standards for watermarking in order to help companies clearly label whether the content is real or AI-generated in order to police, identify and protect against fake AI-generated content. Second, in a move that may impact emerging technologies, various agencies are tasked with studying and reporting on a broad variety of topics related to use of AI and AI-enabled technologies, including generative AI, in the fields of health, education, human services, pharmaceutical delivery and development, and transportation. The result of those studies may impact several emerging technologies, including development of drugs and drug treatments for which biotechnology companies are actively innovating with AI.

Interestingly, and keeping with theme of addressing national security and social risks around the use of AI, the Order somewhat minimizes the threats of IP infringement from generative AI. The Order specifically recognizes that there are "secure and reliable generative AI capabilities, at least for the purposes of experimentation and routine tasks that carry a low risk of impacting Americans' rights." That said, the General Services Administration and Office of Management and Budget, and other relevant agencies, will develop frameworks for prioritizing authorization processes for use of generative AI within the federal government.

Relating to immigration, within 90 days of issuance of the order, the State Department and DHS will streamline processing times and visa applications from noncitizens who wish to come to the US to study or research AI or other critical emerging technologies, and within 120 days of issuance establish new criteria to designate countries and the State Department skills for J-1 immigrants and expand the categories for visa renewals related to AI and emerging technologies. Within 180 days, DHS will evaluate modernizing immigration pathways for experts in AI, including by modernizing the H-1B program and modernizing and expanding the scope of "noncitizens of extraordinary ability."

The Federal Trade Commission (FTC) is tasked with using its authority to ensure fair competition in the AI marketplace and to ensure against harms that may be enabled by the use of AI. For instance, within 180 days of issuance of the order, the Department of Labor (the Labor Department) is required to publish principles and best practices to minimize AI's potential harm to employees' well-being and maximize its benefits, including addressing job standards, job displacement risks and collection and use of employee data. Various federal law enforcement and civil rights agencies are required to report on the effects of AI in the administration of justice. The Department of Health and Human Services (HHS) is also required to publish a plan addressing the use of AI algorithms in the implementations by states and local governments related to public benefits and services. The Labor Department will also need to publish federal contractor guidance regarding the use of AI-based hiring systems. Specifically, the Federal Housing Finance Agency and Consumer Financial Protection Bureau (CFPB) are encouraged to require their regulated entities to use AI to ensure compliance with federal law. The Order directs other agencies to consider, evaluate and revise federal privacy guidelines and standards, including detailed reviews of personally identifiable information obtained by the government or its vendors. The agencies are further directed to protect Americans' privacy by prioritizing federal support for accelerating the development and use of privacy-preserving techniques—including using cutting-edge AI that lets AI systems be trained while preserving the privacy of the training data.

The Order also requires each government agency, within 60 days of issuance, to designate a Chief AI Officer who will be primarily responsible for each such agency's use of AI, promoting AI innovation within the agency, managing AI risks and addressing some of the agency reporting and study requirements in the Order. Certain agencies are required to create AI government boards and will be subject to requirements to develop AI strategies. OMB will issue recommendations to agencies regarding testing of AI (for both accuracy and security), watermarking (to prevent IP theft), establishment of mandatory minimum risk practices, independent evaluation of government contractor's AI claims and documenting and overseeing the use of procured AI by the government.

The Energy Department and NSF will fund the creation of a Research Coordination Network dedicated to advancing privacy research. It is possible that this research will lead to new federal privacy laws, finally aligning the federal government with other privacy regimes to which many U.S. companies are already subject, such as GDPR or California's privacy law regimes.

The Order identifies several hiring initiatives to bring AI expertise into the federal government.

The Order also seeks to find a leading place at the global table for the United States. The State Department, and other agency heads, are tasked with creating strong international frameworks for managing risks and harnessing the benefits of AI. The Commerce Department and State Department are to lead preparations with international allies to develop and implement AI-related consensus standards, cooperation and coordination, including publication of an "AI in Global Development Playbook."

And finally, recognizing that no aspect of the federal government is to be left untouched, the Order creates the White House Artificial Intelligence Council consisting of the Council Chair, the Secretary of State, the Secretary of the Treasury, the Secretary of Defense, the Attorney General, the Secretary of Agriculture, the Secretary of Commerce, the Secretary of Labor, the Secretary of HHS, the Secretary of Housing and Urban Development, the Secretary of Transportation, the Secretary of Energy, the Secretary of Education, the Secretary of Veterans Affairs, the Secretary of Homeland Security, the Administrator of the Small Business Administration, the Administrator of the United States Agency for International Development, the Director of National Intelligence, the Director of NSF, the Director of OMB, the Director of OSTP, the Assistant to the President for National Security Affairs, the Assistant to the President for Economic Policy, the Assistant to the President and Domestic Policy Advisor, the Assistant to the President and Chief of Staff to the Vice President, the Assistant to the President and Director of the Gender Policy Council, the Chairman of the Council of Economic Advisers, the National Cyber Director, the Chairman of the Joint Chiefs of Staff and the heads of such other agencies, independent regulatory agencies and executive offices as the Council Chair may from time to time designate or invite to participate.

The Order is a bold government step in an emerging technology, and the far-reaching consequences of the Order will necessarily be felt for many years as the provisions and directives begin to ramp up. Companies should, however, not ignore the important binding provisions that are scheduled to go into effect 90 days after issuance of the order.

AI is an emerging and rapidly developing area of law. Companies would do well to consider how their business uses or is impacted by AI and take note of the changing legal landscape by getting up-to-date and forward-looking advice from its legal advisors, including considering things like internal AI policies to ensure that companies' use of AI is safe, reliable and effective.

If you have any questions about this article, please contact Daniel H. Shulman at dshulman@vedderprice.com or (312) 609-7530, Sudip K. Mitra at smitra@vedderprice.com or (312) 609-7617 or any other Vedder Price attorney with whom you have worked.