November 7, 2023

The U.S. Executive Order on Artificial Intelligence Is Not Only for Uncle Sam: Key Components of the Executive Order and Potential Impact on the Private Sector

by Joshua S. Ganz , Meghan K. Farmer

On October 30, 2023, the Biden-Harris Administration issued an Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence (the "EO") in an effort to advance American leadership and governance with respect to artificial intelligence ("AI")[1]. Meanwhile, despite hearings and bills introduced by U.S. lawmakers, Congress has yet to make significant progress on comprehensive AI legislation. Given this lack of federal AI legislation, the EO is arguably the biggest development concerning U.S. AI regulation to date.

The EO is intended to ensure the U.S. government can harness the significant benefits of AI while mitigating associated risks. As such, the EO primarily imposes obligations upon U.S. federal agencies, which are summarized in the first part of this alert. However, as discussed in the second part of this alert, the EO will very likely have significant impacts on the creation and use of AI by the private sector and contains valuable takeaways and a roadmap for guidance that will be forthcoming from a variety of federal agencies.

a) The EO's Eight Principles and Associated Key Obligations for U.S. Federal Agencies

Section 2 of the EO sets forth eight principles regarding the development and use of AI. Specific responsibilities associated with these eight principles for U.S. federal agencies are addressed in Sections 4 through 11 of the EO. A high-level overview of these eight principles and key associated responsibilities for U.S. federal agencies include:

1. **Ensuring AI is safe and secure**, which requires testing and evaluation of AI systems to mitigate and address AI-related risks. To further this principle, Section 4.1(a) of the EO tasks the Secretary of Commerce, acting through the Director of the National Institute of Standards and Technology ("NIST"), with establishing guidelines and best practices to promote consensus industry standards for safe, secure, and trustworthy AI systems and establishing safeguards to enable developers of AI to conduct tests to find flaws and vulnerabilities in AI systems. Meanwhile, Section 4.5(c) of the EO requires the Secretary of Commerce to issue guidance to agencies for labeling AI generated content (such as "watermarking") so Americans can determine when content is authentic.

2. **Promoting responsible innovation, competition, and collaboration**, which requires investment in AI while simultaneously tracking intellectual property questions and promoting a fair, open, and competitive ecosystem and marketplace for AI. To further the intellectual property component of this principle, under Section 5.2(c) of the EO, the Under Secretary of Commerce for Intellectual Property and Director of the United States Patent and Trademark Office ("USPTO") are tasked with issuing guidance to USPTO patent examiners and applicants addressing inventorship and the use of AI and other considerations at the intersection of AI and intellectual property.
3. **Supporting American workers**, which requires improving workers' lives and preventing AI from being deployed in the workplace in a manner that undermines rights, worsens job quality, encourages undue worker surveillance, lessens market competition, introduces new health and safety risk, or causes harmful labor-force disruptions. Under Section 6(b) of the EO, the Secretary of Labor is charged with developing and publishing principles and best practices for employers that could be used to mitigate AI's potential harms to employees' well-being and maximize its potential benefits, which include specific steps for employers to take regarding AI.
4. **Implementing AI policies consistent with advancing equity and civil rights**, which includes complying with existing federal laws and promoting robust technical evaluations, careful oversight, engagement with affected communities, and regulation. Section 7.1 of the EO requires the Attorney General to coordinate with and support agencies in their implementation and enforcement of existing federal laws to address civil rights and civil liberties violations and discrimination related to AI.
5. **Protecting the interests of Americans who use, interact with, or purchase AI and AI-enabled products**, including encouraging the enforcement of existing consumer protection laws and principles and enacting appropriate safeguards against fraud, unintended bias, discrimination, infringements on privacy, and other harms from AI. Section 8(a) of the EO encourages independent regulatory agencies to address consumer risks that may arise from AI, while Sections 8(b) – (e) of the EO provide for specific actions for U.S. federal agencies in the healthcare, public health, human services, transportation, education, and communication sectors.
6. **Protecting Americans' privacy and civil liberties** by ensuring the collection, use, and retention of data is lawful, secure, and mitigates privacy and confidentiality risks. Section 9(a) of the EO requires the Director of Office of Management and Budget to: (i) identify commercially available information, including personally identifiable information procured by government agencies from third parties; and (ii) establish standards and procedures aimed at mitigating privacy and confidentiality risk regarding the collection and processing of such information. Section 9(b) of the EO requires the Secretary of Commerce to create guidelines for agencies to evaluate the efficacy of "differential-privacy-guarantee" protections including for AI (discussed further below). Last, Section 9(c) of the EO focuses on the U.S. government's research, development, and implementation of "privacy-enhancing technologies" or "PETs", defined in relevant part as "… any software or hardware solution, technical process, technique, or other technological means of mitigating privacy risks arising from data processing…" in connection with AI.
7. **Managing the risks and ensuring effective federal government use of AI**, including through public service-oriented AI professionals, and ensuring training for the government's workforce. Section 10.1(a) requires the Director of Office of Management and Budget to convene and chair an interagency counsel to coordinate the development and use of AI in agencies' program and operations.

8. **Leading the way to global societal, economic, and technological progress** through engaging with internationally allies and partners to develop a framework for AI and promote responsible AI safety and security principles with other nations. Section 11(ii) of the EO requires the Secretary of State to lead efforts to establish a strong international framework along with common principles for managing the risks and harnessing the benefits of AI.

b) Applicability of the EO on the Private Sector and Practical Suggestions for AI Governance

Although primarily directed to U.S. federal agencies and their rulemaking power, the EO may also impose requirements on organizations in the private sector. Section 4.2(a) of the EO mandates the Secretary of Commerce to require companies with an intent to develop a potential "dual-use foundation model"[2] to provide the federal government with ongoing reports and also to require companies or individuals to report on their acquisition, development, and possession of a potential "large-scale computing cluster".[3] In addition, Section 4.2(c)(i) requires the Secretary of Commerce to implement reporting requirements for United States Infrastructure as a Service Providers when a foreign person transacts with such provider to retain a large AI model with potential capabilities that could be used in a malicious cyber-enabled activity.

Significantly, the guidance to be issued under the EO can strongly inform best practices for the private sector with respect to the creation, use, and deployment of AI services and mitigating risk associated with AI. Key focus areas that will likely be broadly applicable to many companies include:

1. **NIST Guidelines and Best Practices.** As mentioned above, the NIST will be charged with developing guidelines and best practices for deploying safe, secure, and trustworthy AI systems and to develop a companion resource to the existing NIST AI Risk Management Framework within 270 days of the date of the EO. So, while waiting on forthcoming guidelines and best practices, organizations may want to consider evaluating their existing AI risk management framework against the current NIST AI Risk Management Framework. This will assist organizations in developing a baseline for their AI risk management program while additional guidance is forthcoming.
2. **Principles and Best Practices with Respect to Employees and Internal AI Use.** The principles and best practices to be published by the Secretary of Labor discussed above will be forthcoming within 180 days of the date of the EO. These principles and best practices will likely inform organizations whether and how to best utilize AI with respect to their own employees and within their own organization internally. Meanwhile, as discussed in part in a previous article, organizations should continue to carefully select their AI tools and to implement internal governance surrounding their AI tools. Practically, companies should develop an internal AI policy governing internal uses, safeguards, and oversight of AI and continue to refine the policy as the mandated rulemaking unfolds.
3. **Data Privacy and Security Guidance is Forthcoming.** Those charged with leading or supporting an organization's data privacy and security program will likely want to review the Secretary of Commerce's guidelines to evaluate the efficacy of "differential-privacy-guarantee" protections under Section 9(b) of the EO discussed above, which will finalized within 365 days of the EO. Section 3(j) of the EO defines "differential-privacy guarantee" as "…protections that allow information about a group to be shared while provably limiting the improper access, use, or disclosure of personal information about particular entities."

At a minimum, these forthcoming guidelines must include a description of the significant factors that bear on differential-privacy safeguards and common risks to realizing differential privacy practice. These guidelines may not only be helpful with respect to AI but may also be helpful to refine an organization's data privacy and security program generally.

4. **Other AI Governance Considerations.** Last, given the EO's primary application to the U.S. government, it is important to note that congressional action to institute holistic federal AI legislation for the private sector is likely necessary. Without federal AI legislation, similar to data privacy law in the U.S., the U.S. could end up with a patchwork of AI legislation at the state level. While awaiting further developments on AI legislation, organizations using and offering AI products should still take action to address the impacts of AI. What to prioritize and how to do this effectively can vary depending on how an organization uses AI or offers AI products. For example, if an organization's third-party vendors are utilizing AI to provide products or perform services for an organization (e.g., a marketing agency is utilizing AI to create marketing materials), contracts with such vendors should include AI-related terms. If your organization is offering products or service that leverage AI, customer agreements or terms should contain key provisions that mitigate the risks associated with providing those products (e.g., disclaiming warranties regarding the outputs that AI produces). In addition, it is important to consider existing principles under consumer protection and data privacy laws that are applicable to AI, such as providing adequate transparency to consumers regarding AI practices and lawful and responsible data collection practices in the development of an AI product.

Footnotes

---

[1] In Section 3(b) of the EO, "artificial intelligence" or "AI" is defined broadly as "a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments," which makes the EO applicable to a wide range of AI systems.

[2] Under Section 3(k) of the EO, "dual-use foundation model" means "…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 easily be 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: (i) substantially lowering the barrier of entry for non-experts to design, synthesize, acquire, or use chemical,, biological, radiological, or nuclear (CBRN) weapons; (ii) enabling powerful offensive cyber operations through automated vulnerability discovery and exploitation against a wide range of potential targets of cyber attacks; or (iii) permitting the evasion of human control or oversight through means of deception or obfuscation." The EO clarifies in the definition that AI 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.

[3] "Large-scale computing cluster" is undefined under the EO. However, under Section 4.2(b) of the EO, it clarifies that technical conditions for "dual-use foundation models" and "large-scale computing clusters" will be forthcoming. However, Section 4.2(b) of the EO also clarifies until such technical conditions are defined, the Secretary of Commerce must require compliance with these reporting requirements for: (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.