LATHAM&WATKINS

Client Alert Commentary

Latham & Watkins Artificial Intelligence Taskforce November 2, 2023 | Number 3187

President Biden's Executive Order on Artificial Intelligence — Initial Analysis of Private Sector Implications

The Order marks an ambitious effort to stand up a whole-of-government approach to encouraging the benefits and managing the risks of artificial intelligence, with many of its most significant private-sector implications announced but not yet in place.

On October 30, 2023, President Biden issued a sweeping Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence (with an accompanying Fact Sheet). The Executive Order appears designed to balance the goals of fostering innovation and remaining globally competitive, with intense pressure to secure and regulate AI in the United States. It is structured around eight "guiding principles and priorities":

- Ensuring the Safety and Security of AI Technology: Government agencies must develop AI safety standards, and developers of large AI models and those with large-scale computing clusters that meet certain technical requirements must share significant information with the government.
- Promoting Innovation and Competition: The United States must catalyze AI research, promote a fair and competitive AI ecosystem, and expand the pool of individuals with AI expertise.
- Supporting Workers: Federal agencies associated with labor and the workforce, such as the Department of Labor, must develop principles and best practices to mitigate the harms and maximize the benefits of AI for workers.
- Advancing Equity and Civil Rights: Federal agencies across the government are directed to develop new guidance, plans, and other measures within the scope of their authorities to combat the potential for discrimination and other harms that AI may exacerbate.
- Protecting Consumers, Patients, Passengers, and Students: Various agencies must ensure that AI is used safely and responsibly in healthcare, education, and transportation, and that consumers are protected from fraud, discrimination, and privacy risks related to AI.
- Protecting Privacy: Federal agencies, such as the Department of Commerce and OMB, must study and strengthen privacy-enhancing technologies (PETs) and other privacy measures for the use of data in AI, particularly for the government's use of AI systems.

- Advancing Federal Government Use of AI: The federal government must develop guidance for agencies' use of AI, acquire AI products and services, and hire AI professionals.

- Strengthening American Leadership Abroad: The Biden administration will work with other nations to support the safe, secure, and trustworthy deployment and use of AI worldwide.

Many of the Order's provisions serve to direct federal agencies to study the effects of AI within their regulatory purview and provide policy recommendations to the President. The Order also encourages agencies — including the FTC, CFPB, and HHS specifically — to enforce existing consumer protection laws to address risks that arise with use of AI, and to engage in further rulemaking specific to AI. The Order also creates a new White House AI Council, with representatives of 29 federal divisions and agencies to coordinate on implementation progress and efforts.

For companies and individuals, there are three overarching points:

First, under the principle of "Ensuring the Safety and Security of AI Technology," the White House is viewing large AI models and the large computing clusters capable of use in their development and training as matters of national security. The Order invokes the Defense Production Act (DPA) (50 U.S.C. § 55) to impose requirements on companies developing (or intending to develop) AI models that meet certain technical processing thresholds, and on entities and individuals that possess certain large-scale computing clusters. Under this authority, developers of large AI models must implement sufficient physical and cybersecurity efforts to protect models and model weights from foreign infiltration, and must comply with detailed reporting requirements to the federal government. In addition, entities and individuals that possess a potential large-scale computing cluster will need to report to the government the existence and location of each computer cluster and the amount of total computing power available in each cluster. These reporting requirements are the most specific new rules that come out of the Executive Order. However, as of today, they apply only to companies whose models meet the definition of "dual-use foundation models"[1] and meet certain technical specifications,[2] or that possess or offer computing centers meeting certain thresholds.[3] Going forward, the Secretary of Commerce will define and update the set of technical conditions for models and computing clusters subject to these reporting requirements.

Second, the Order calls on federal regulators to regulate aggressively the use of AI in their respective areas of enforcement, such as consumer protection, civil rights, education, financial opportunities, transportation, and healthcare. The potential impact to business is much broader than the national security-based requirements for developing models or possessing computing clusters; any business using AI must protect against sector-specific risks from their use of AI. Key principles underlying the regulatory call to action are "Advancing Equity and Civil Rights," "Protecting Consumers, Patients, Passengers and Students," and "Protecting Privacy."

Third, the requirements for AI will remain dynamic. The Order calls specifically for government studies as well as future regulations across a wide swath of the federal government. Notable provisions include requirements for:

- The National Institute of Standards and Technology to develop — as the Fact Sheet describes the requirement — "rigorous standards" for "extensive red-team safety testing to ensure safety before public release" (Section 4.1(a)(ii))

- The Department of Commerce to draft guidance for detecting and authenticating AI content (Section 4.5(a))

- The FTC to use its rulemaking authority "to ensure fair competition in the AI marketplace and to ensure that consumers and workers are protected from harms that may be enabled by the use of AI" (Section 5.3(a))

- The Secretary of Commerce to issue regulations concerning reporting requirements related to the use of certain AI models by foreign persons (Section 4.2(c), (d))

The Executive Order directs relevant federal agencies to take action under relatively short timelines, generally ranging from 45 to 365 days. Below we identify and describe the provisions of the Executive Order that will apply directly to businesses and provide initial analysis of the Executive Order.

Potential Impact on Businesses

National Security and Invocation of the DPA

An issue of national security: Under Section 4.2(a), "Ensuring Safe and Reliable AI," certain companies will be required to share certain information with the federal government under the authority of the Defense Production Act (DPA) (50 U.S.C. § 55). Passed in 1950 as part of the Cold War mobilization effort, presidents have used the DPA in a variety of circumstances. While the DPA is most often used to procure materials critical to national security, it has also been employed in such wide-ranging circumstances as requiring telecommunications companies to provide information about their networks, counter spying, and mandating that General Motors make ventilators during the COVID-19 pandemic.

As detailed below, the Order invokes the DPA to mandate that certain companies share particular information with the government, including information related to model training, the ownership and possession of model weights and physical and cybersecurity measures taken to protect such weights, red-team testing results, and location and power of computing clusters. Notably, the White House has relied on emergency powers under the DPA to impose these requirements, signaling that the White House views large AI models as a national security issue. In addition to specific reporting requirements, developers working on such models generally will need to implement security controls of the highest degree to protect against foreign infiltration.[4]

Covered companies: Section 4.2(a) applies to:

a) companies "developing or demonstrating an intent to develop potential dual-use foundation models"[5] that meet certain thresholds, and

b) "companies, individuals, or other organizations or entities that acquire, develop, or possess a potential large-scale computing cluster" used to train AI.

Notably, many current AI models and computer clusters may not meet the thresholds of these requirements. The Order sets forth initial thresholds for models and computing clusters, and gives the Department of Commerce the authority to define the kinds of models and computer clusters that are subject to these reporting requirements (Order section 4.2(b)). Until the Secretary of Commerce defines those conditions, the initial thresholds are:

- *Models*: 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

- *Computing clusters*: Any computing cluster that has a set of machines physically co-located in a single datacenter, is transitively connected by data center networking of over 100 Gbit/s, and has a theoretical maximum computing capacity of 10^{20} integer or floating-point operations per second for training AI.[6]

Requirements: Companies whose models meet the above requirements for "potential dual-use foundation models" must, starting within 90 days, "provide the Federal Government, on an ongoing basis, with information, reports, or records regarding the following":

- Any ongoing or planned activities related to training, developing, or producing dual-use foundation models, including the physical and cybersecurity protections taken to protect the training process against sophisticated threats.

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

- The results of any dual-use foundation model's red-team testing based on guidance developed by NIST pursuant to the Order, and a description of any "associated measures" the company takes to meet safety objectives, such as mitigations to "improve performance on red-team tests and to strengthen overall model security".[7] Note that until NIST guidance is issued, the reporting requirement applies to "any red-team testing" conducted on certain topics.[8]

Companies that possess "a potential large-scale computing cluster" must report any such possession, including information about the existence and location of the clusters and the amount of available computing power in each cluster.

The additional requirement to produce red-teaming efforts is likely to be the most consequential. The red-teaming reports should cover a listed set of potential risks and are subject to further guidance from NIST. Therefore, the Order will require AI developers to reassess their approach to developing their models. For instance, if the process of red-teaming involves competitive or otherwise sensitive information, disclosure to the government could deprive that information of existing protections. Going forward, developers will need to conduct red-teaming with an eye on the Order requirements and be mindful that such reports will need to be produced. Companies may want to develop a specific "red team" so they can clearly separate that team's work from the product team or others that may also be testing the model in development.

Focus on Regulatory Protection Against AI Risks

The Order is also a call to action for the federal government to dedicate substantial energy and resources to addressing potential risks associated with AI. The Order sets forth a clear mandate for rulemaking and enforcement. Most applicable to businesses, the Order specifically calls for regulation and enforcement by the following agencies to address the following risks:

- Federal Trade Commission (FTC): The Order encourages the FTC to use its existing powers and rulemaking authority "to ensure fair competition in the AI marketplace and to ensure that consumers and workers are protected from harms that may be enabled by the use of AI" (Section 5.3(a)).

- The Department of Health and Human Services (HHS): The Order mandates that within 90 days, HHS must establish an AI Task Force that shall, within 365 days, develop a strategic plan — potentially including regulatory action — to deploy AI in the health and human services sector

(including research and discovery, drug and device safety, healthcare delivery and financing, and public health) (section 8(b)).

- Consumer Financial Protection Bureau (CFPB): The Order encourages the CFPB to use its existing authority to require the entities it regulates to use appropriate AI methodologies to ensure compliance with federal law, including to minimize bias or disparities for protected groups (section 7.3(b)).

- Secretary of Education: The Order requires the Secretary of Education to "develop resources, policies, and guidance" regarding AI that "address safe, responsible, and nondiscriminatory uses of AI in education," including an "AI toolkit" for education leaders (section 8(d)).

- Secretary of Transportation: The Order mandates that the Secretary of Transportation take various actions in relation to AI, including exploring transportation-related opportunities and challenges, and assessing the need for information, technical assistance, and guidance related to AI (section 8(c)).

- Secretary of Housing and Urban Development (HUD): The Order identifies a concern with AI discrimination and biases in "access to housing and in other real estate-related transactions," and requests that HUD issue guidance on how AI may violate federal law and best practices companies can employ to avoid such violations (Section 7.3(b), (c)).

Given this sector-specific approach and the call for regulatory attention to these issues, any business developing or implementing AI models will need to understand whether their use of AI raises any sector-specific risks and whether they have undertaken sufficient action to protect against those risks.

Additional Forthcoming Noteworthy Regulations and Rules

The Order requires many different federal agencies to study AI and make policy recommendations, or promulgate additional regulations and policies. We do not cover all of these requests in this Alert, as topics range from defense to veterans, clean energy to healthcare, intellectual property to communications networks, criminal justice to scientific research, to accessibility and so on.

Some of the more impactful regulations for businesses include requests to federal agencies to develop guidance to help detect and label synthetic content generated by AI (section 4.5(a)), and to issue recommendations on copyright and IP-related risks (section 5.2(c), (d)). In addition, some of the more concrete requests for rules and regulations that may impact businesses include:

Use of Certain Models by Foreign Persons: Section 4.2(c) of the Executive Order requires the Secretary of Commerce to propose regulations governing the use of certain models by foreign persons within 90 days:

Covered companies: The regulations apply to "Information as a Service Providers" or "IaaS Providers" when a foreign person transacts with an IaaS Provider to train a large AI model with "potential capabilities to be used in a malicious cyber-enabled activity (a 'training run')."[9] The Secretary of Commerce must determine the technical conditions that would allow potential for an AI model to be used in malicious cyber-enabled activity.

Until then, a model shall be considered to have potential for such capabilities if it requires a quantity of computing power greater than 10^{26} integer or floating-point operations and is trained on a computing cluster with 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 compute capacity of 10^{20} integer or floating-point operations per second for training AI.

Requirements: The regulations must require IaaS Providers to identify any foreign person who transacts with them to train a large AI model, by reporting to the government the identity of the foreign person, the existence of the training run, and additional information to be determined by regulation. IaaS Providers must also prohibit any foreign reseller of the product from providing the products unless the foreign reseller submits a report to the Provider, which the Provider must submit to the Secretary of Commerce, detailing each instance in which a foreign person transacts with the foreign reseller to use the United States IaaS Product to conduct a training run. These rules could apply, for example, to large cloud providers that rent computing space for AI purposes to foreign customers.

Verification of Foreign Persons Using Certain AI Models. Section 4.2(d) mandates that within 180 days the Secretary of Commerce propose regulations obligating domestic IaaS Providers to require foreign resellers of IaaS Products to verify the identity of foreign persons who obtain an IaaS account.

Covered companies: IaaS Providers

Requirements: The regulations shall:

- Provide the minimum standards an IaaS Provider must require of foreign resellers to verify the identity of an individual who creates an account with the foreign reseller; and
- Establish regulations that foreign resellers of IaaS Products must follow if they allow AI models using domestic IaaS Providers to be used by foreign individuals to conduct training runs, including identifying each instance in which a foreign person conducts such a training run

New Standards: Section 4.1 of the Executive Order mandates that the National Institute of Standards and Technology establish guidelines and best practices within 270 days, "with the aim of promoting consensus industry standards for developing and deploying safe, secure, and trustworthy AI systems." NIST shall fulfill this obligation by:

- Developing a companion resource to the AI Risk Management Framework, NIST AI 100-1, for generative AI;
- Developing a companion resource to the Secure Software Development Framework to incorporate secure development practices for generative AI and for dual-use foundation models; and
- Launching an initiative to create guidance and benchmarks for evaluating and auditing AI capabilities, with a focus on capabilities through which AI could cause harm, such as in cybersecurity and biosecurity.

NIST must establish guidelines to enable AI developers "to conduct red-teaming tests to enable deployment of safe, secure, and trustworthy systems," including:

- Developing guidelines to assess and manage the safety, security, and trustworthiness of dual-use foundation models; and
- Developing and helping to ensure the availability of testing environments, such as testbeds, to support the development of safe, secure, and trustworthy AI technologies.

Widely Available Model Weights: After gathering input from the private sector and other stakeholders via a public consultation process, the Secretary of Commerce shall submit a report to the President on the potential implications of dual-use foundation models for which the model weights are widely available, as well as policy and regulatory recommendations pertaining to those models.[10] Businesses may want to provide input into this process, given that the Secretary is tasked expressly with soliciting input from the private sector.

Practical Takeaways

The Executive Order is a significant step towards furthering the governance of AI in the United States and beyond. And the timing of the Executive Order's issuance is notable. It was issued contemporaneously with the Hiroshima AI Process Comprehensive Policy Framework—a set of international guiding principles on AI and a voluntary code of conduct agreed by the leaders of the G7—also published on October 30, 2023 (G7 Guiding Principles). It was also published just days before the AI Safety Summit, hosted by the UK Government, at which key figures in the AI industry were to gather how to best manage the risks posed by the most recent advancements in AI. There are notable similarities between the Executive Order and the G7 Guiding Principles which support a risk and principle-based approach to the regulation of AI. If the actions the Executive Order directs are executed and the broad range of regulations it contemplates are adopted, businesses developing, supporting, or using AI and large computing clusters will be subject to increased regulations and risks. The Executive Order, however, could also present significant opportunities for technology companies and other companies to support their customers in implementing these requirements. We will watch closely as the next steps develop.

If you have questions about this Client Alert, please contact one of the authors listed below or the Latham lawyer with whom you normally consult:

Michael H. Rubin
michael.rubin@lw.com
+1.415.395.8154
San Francisco

Andrew Gass
andrew.gass@lw.com
+1.415.395.8806
San Francisco

Ghaith Mahmood
ghaith.mahmood@lw.com
+1.213.891.8375
Los Angeles

You Might Also Be Interested In

SEC Proposes New Rules Targeting Use Predictive Data Analytics by Investment Advisers and Broker-Dealers

IP and Privacy 101: A Crash Course for Busy Entrepreneurs

ESG, Technology, and AI: The Next Evolution

Shooting for the Moon: The Evolution of Key AI/ML Regulations Governing Certain Health Care Products and Services

Endnotes

[1] "The term "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 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:

(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 cyberattacks; or

(iii) permitting the evasion of human control or oversight through means of deception or obfuscation.

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."

[2] The thresholds will be updated by the Secretary of Commerce but until then are set at the following: "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." Order Section 4.2(b).

[3] Until the thresholds are updated by the Secretary of Commerce they are set at the following: "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." Order section 4.2(b).

[4] Administration officials have confirmed the intended application of the Executive Order, with White House Chief of Staff Jeff Zients stating, "At the end of the day, the companies can't grade their own homework here. So we've set the new standards on how we work with the private sector on AI, and those are standards that we're going to make sure the private companies live up to." Similarly, in a briefing to reporters, an anonymous administration official noted that the Executive Order applies to companies' most powerful AI systems — regardless of whether the companies work with the federal government.

[5] "The term "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 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:

(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 cyberattacks; or

(iii) permitting the evasion of human control or oversight through means of deception or obfuscation.

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."

[6] *See* Order Section 4.2(b).

[7] The Executive Order defines "red-teaming" as "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. Artificial Intelligence 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."

[8] 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; the possibility for self-replication or propagation; and associated measures to meet safety objectives." Section 4.2(a)(i)(C).

[9] The Executive Order defines "Infrastructure as a Service Provider" as any United States entity that offers Infrastructure as a Service Product. "Infrastructure as a Service Product" means any product or service offered to a consumer, including complimentary or "trial" offerings, that provides processing, storage, networks, or other fundamental computing resources, and with which the consumer is able to deploy and run software that is not predefined, including operating systems and applications. The consumer typically does not manage or control most of the underlying hardware but has control over the operating systems, storage, and any deployed applications. The term is inclusive of "managed" products or services, in which the provider is responsible for some aspects of system configuration or maintenance, and "unmanaged" products or services, in which the provider is only responsible for ensuring that the product is available to the consumer. The term is also inclusive of "virtualized" products and services, in which the computing resources of a physical machine are split between virtualized computers accessible over the Internet (e.g., "virtual private servers"), and "dedicated" products or services in which the total computing resources of a physical machine are provided to a single person (e.g., "bare-metal" servers).

[10] *See* Order Section 4.6.