

CLIENT ALERT | June 3, 2026

President Trump Signs Executive Order Establishing AI Cybersecurity and Frontier Model Framework

The Trump administration has issued a light-touch executive order focused on cyber risks stemming from frontier AI.

Key Points:

- The Order directs federal agencies to strengthen cybersecurity across government systems and critical infrastructure, including through an AI cybersecurity clearinghouse and expanded access to AI-enabled defensive tools for state and local authorities.
- Federal agencies will design a voluntary framework by August 1, 2026, for developers of frontier AI models to engage with the federal government prior to model release.
- The Attorney General is directed to prioritize enforcement of existing federal criminal statutes against anyone who uses AI to illegally access or damage a computer without authorization, or who employs AI agents to unlawfully access data for use in a criminal purpose.

On June 2, 2026, President Trump signed an executive order (the Order) titled “Promoting Advanced Artificial Intelligence Innovation and Security,” which seeks to balance the Trump administration’s long-standing goal of promoting AI innovation through a minimally burdensome regulatory framework with the need to modernize and protect key systems against emerging AI-related threats.

The Order addresses three principal areas: (1) strengthening cybersecurity across federal systems and critical infrastructure; (2) creating a voluntary pre-release engagement framework for developers of frontier AI systems; and (3) directing the Attorney General to prioritize enforcement of existing criminal statutes against AI-enabled cybercrimes. Throughout, the Order emphasizes voluntary collaboration with the AI industry rather than implementing rigid, mandatory regulations.

This Client Alert summarizes the Order’s key provisions and outlines practical steps that companies developing, deploying, or incorporating AI into their business operations should consider.

Overview of the Order

The Order is organized around three core components: (I) upgrading federal and critical infrastructure cybersecurity; (II) establishing a voluntary framework for pre-release engagement on frontier AI models; and (III) prioritizing criminal enforcement against AI-enabled cybercrimes.

I. Upgrading American Systems for Advanced AI

The Order directs an effort to strengthen federal cybersecurity and expand access to AI-enabled defensive tools. Within 30 days (i.e., by July 2, 2026), the Committee on National Security Systems and the Secretary of War must each prioritize the cyber defense of their respective information systems. Likewise, the Secretary of Homeland Security, through the Cybersecurity and Infrastructure Security Agency (CISA), must issue Binding Operational Directives to expedite the cyber defense of civilian federal systems; establish or expand federal programs that enhance AI-enabled defensive tools; and facilitate access to cybersecurity tools for agencies, state and local authorities, and critical infrastructure operators such as rural hospitals, community banks, and local utilities.

Also within 30 days, the Secretary of the Treasury must form an “AI cybersecurity clearinghouse” in voluntary collaboration with the AI industry and critical infrastructure operators to coordinate vulnerability scanning, discover and validate vulnerabilities, and coordinate and prioritize remediation and patch distribution. The Director of the Office of Management and Budget must determine whether federal grant programs have funding available for applicants developing advanced AI vulnerability detection. Within 60 days (i.e., by August 1, 2026), the Director of the Office of Personnel Management must expand the hiring pathways for US Tech Force information cybersecurity specialists, which could have potential downstream effects on private-sector labor markets.

II. Secure Frontier Model Deployment

The Order directs federal officials to develop, within 60 days, both a classified benchmarking process for assessing the advanced cyber capabilities of AI models and a voluntary pre-release engagement channel between developers and the federal government. Critically, the framework is expressly voluntary for developers and does not constitute a mandatory licensing or pre-clearance regime.

The framework will apply to “covered frontier models,” although the Order notably leaves that term undefined. Rather, within 60 days, the Secretary of the Treasury, the Secretary of War (through the National Security Agency), and the Secretary of Homeland Security (through CISA), in consultation with other senior officials and agencies as appropriate, must develop a classified benchmarking process to assess AI models’ advanced cyber capabilities and determine the threshold at which AI models will be designated as a “covered frontier model” under the Order. Because the benchmarking process is classified, AI developers building models with significant cybersecurity-relevant capabilities should anticipate the need to engage with the government to understand whether their models meet this designation.

The Order directs that the framework should enable developers to provide government access to their covered frontier models for up to 30 days before the developer's planned release date for the models to "other trusted partners" (also an undefined term), subject to appropriate confidentiality, cybersecurity, insider-risk, and intellectual-property requirements covering protection, use, and nondisclosure. The framework should also enable developers to collaborate with the federal government to select trusted partners that will have early access to covered frontier models in order to promote secure innovation and strengthen the cybersecurity of critical infrastructure. Notably, the envisioned 30-day window would run prior to the planned release to these "trusted partners" — not to the general public — meaning government access would occur at an earlier stage in the distribution chain than a pre-public-release framework would require. The specific terms of IP and confidentiality protections, as well as the criteria for selecting trusted partners, remain to be developed.

Underscoring the voluntary nature of participation in this framework, the Order expressly states that it does not create any mandatory governmental licensing, pre-clearance, or permitting requirement for the development, publication, release, or distribution of new AI models, including frontier models.

III. Protection Against Criminal Actors

Finally, the Order directs the Attorney General to prioritize the enforcement of existing federal criminal statutes¹ against anyone who utilizes AI to illegally access or damage a computer without authorization, or who utilizes AI while engaged in such illegal access to further any other crime. This includes breaching any public or private information technology system, or employing AI agents to unlawfully access data or information that is subsequently used for a criminal or unlawful purpose.

Strategic Implications for AI-Focused Enterprises

The Order does not impose any mandatory obligations on AI developers, nor does it define most of the key terms and concepts that will ultimately determine its practical impact.

For most companies, the near-term posture should be to monitor the agency guidance expected within the next 30 to 60 days and assess how those developments affect their operations as the regulatory picture comes more clearly into focus. That said, the Order creates several practical considerations that companies can begin evaluating now.

- **Scope of "covered frontier models."** A threshold question for developers of large-scale AI models will be whether their models fall within the scope of the Order's "covered frontier model" designation. The benchmarking process is classified, meaning the precise capability threshold may never be publicly disclosed — in fact, the Order calls for the framework to allow developers to "engage the Federal Government to determine whether model(s) under development meet the designation of 'covered frontier model.'" As such, companies developing large-scale AI models with advanced cyber-relevant capabilities should consider engaging directly with federal agencies shortly after the benchmarking process concludes to understand next steps.

Although the Order does not define “covered frontier model,” other recently enacted legislation offers useful reference points for scoping a potential federal definition. California’s Transparency in Frontier Artificial Intelligence Act, New York’s RAISE Act, and Illinois’ just-passed SB 315 all define a “frontier model” as a foundation model trained using more than 10^{26} integer or floating-point operations (FLOPS). Similarly, the EU AI Act presumes that general-purpose AI models trained with more than 10^{25} FLOPs pose “systemic risk,” triggering additional obligations around model evaluation, incident reporting, and cybersecurity.

While the federal government’s ultimate definition of “covered frontier model” may be shaped by different considerations — particularly advanced cyber capabilities rather than raw compute — these existing thresholds nonetheless provide a useful frame of reference. Given the Trump administration’s stated emphasis on maintaining a lean regulatory framework and avoiding overly burdensome constraints on AI development, the federal definition is unlikely to capture a materially broader universe of models than these existing legislative benchmarks.

- **Participation in the voluntary pre-release framework.** The voluntary framework for pre-release government engagement is due to be finalized by August 1, 2026. While the Order generally contemplates a framework that would guide developers on how to provide the federal government with access to covered frontier models for up to 30 days before release to “trusted partners,” the specific mechanics, scope, and conditions of that framework have not yet been designed, making it difficult to assess the potential benefits and drawbacks of voluntary participation at this point.

For example, the process and criteria for how the government and developers will collaborate to select “trusted partners” eligible for early access to covered frontier models have not been articulated, although the Order states that the purpose of such collaboration is to promote secure innovation and strengthen the cybersecurity of critical infrastructure. Likewise, the Order requires that developers who participate in the voluntary framework be afforded “appropriate confidentiality, cybersecurity, insider-risk, and intellectual-property protection, use, and nondisclosure requirements.” The specifics of those limitations on government access to model weights or other proprietary information will be a material consideration for many developers in deciding whether or not to participate.

- **Opportunities and uncertainties of the AI cybersecurity clearinghouse.** The Treasury-led AI cybersecurity clearinghouse could create meaningful benefits for participants. Participation in the clearinghouse could give AI developers and critical infrastructure operators early access to coordinated vulnerability intelligence, such as government-identified software vulnerabilities, validated threat assessments, and prioritized remediation and patch distribution. Participation could also position companies to shape the clearinghouse’s norms and practices during its formative period.

However, the clearinghouse's governance structure, data-sharing protocols, and liability protections for participants remain undefined. Companies should consider how vulnerability information shared with the government could be used, stored, or disclosed, and whether sharing details about internal vulnerabilities could create unique regulatory, litigation, or reputational risk. For instance, the Order does not address whether information submitted to the clearinghouse would be shielded from public disclosure under laws like the Freedom of Information Act or barred from use in any future enforcement actions.

- **Potential heightened enforcement risk for agentic deployers.** The Order directs the Attorney General to prioritize enforcement under federal criminal statutes against anyone who “employ[s] AI agents to unlawfully access data or information that is subsequently used for a criminal or unlawful purpose.” Companies deploying or preparing to deploy AI agents should carefully assess their potential exposure under the statutes cited in the Order, particularly where agents interact with third-party systems or data.

Moreover, as AI agents become capable of executing purchases and initiating transactions autonomously, companies should ensure robust access controls, permission scoping, and audit logging are in place to prevent agents from creating potential liability issues by completing actions that exceed the intended scope of their authorization.

Contacts

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

Sy Damle

sy.damle@lw.com
+1.202.637.3332
+1.212.906.1659
Washington, D.C. / New York

Fiona Maclean

fiona.maclean@lw.com
+44.20.7710.1822
London

This publication is produced by Latham & Watkins as a news reporting service to clients and other friends. The information contained in this publication should not be construed as legal advice. Should further analysis or explanation of the subject matter be required, please contact the lawyer with whom you normally consult. The invitation to contact is not a solicitation for legal work under the laws of any jurisdiction in which Latham lawyers are not authorized to practice. See our [Attorney Advertising and Terms of Use](#).

Endnotes

-
- ¹ The Order specifically cites 18 U.S.C. § 1028 (fraud and related activity in connection with identification documents), 18 U.S.C. § 1030 (the Computer Fraud and Abuse Act), 18 U.S.C. § 1343 (wire fraud), as well as “all other applicable federal criminal laws.”