UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF COLUMBIA

THE PROTECT DEMOCRACY PROJECT, INC., *Plaintiff*, v. OFFICE OF THE NATIONAL CYBER DIRECTOR, et al., *Defendants*.	Civil Action No.

DECLARATION OF JUSTIN HENDRIX IN SUPPORT OF PLAINTIFF'S MOTION FOR A PRELIMINARY INJUNCTION

I, Justin Hendrix, declare as follows pursuant to 28 U.S.C. § 1746:

I. BACKGROUND

1. I am the Chief Executive Officer and Editor of Tech Policy Press, a nonprofit media organization that publishes journalism, analysis and perspective on issues at the intersection of technology and public policy, including artificial intelligence governance. I make this declaration based on my personal knowledge and my role overseeing Tech Policy Press's editorial operations. If called as a witness, I could and would testify competently to the matters stated herein.

2. Tech Policy Press publishes original reporting, expert analysis, a podcast, and daily and periodic digests covering technology policy developments. Its readership includes congressional and executive-branch staff, state officials, academic researchers, journalists at other outlets, civil-society organizations, and members of the public who follow technology

governance closely. We reach millions of unique visitors each year on our website and social media channels, serve tens of thousands of newsletter subscribers, and publish one of the most popular podcasts on tech policy matters. Tech Policy Press is a 501(c)3 that receives funding from major foundations and individual donors.

3. Before founding Tech Policy Press, I was the Executive Editor of NYC Media Lab, a consortium of universities and media companies, and before that I worked for more than a decade at The Economist. I am also an adjunct professor at the New York University Tandon School of Engineering and a lecturer at Cornell Tech. In this declaration, I take no position on what substantive AI policy the federal government should adopt or on the conduct of any company. I describe how the secrecy of the federal frameworks at issue in this case concretely impairs journalism about the government's regulation of artificial intelligence, and why the value of the requested records to that journalism and its audience is time-bound in ways that ordinary processing timelines would undercut.

II. TECH POLICY PRESS'S COVERAGE OF THIS SUBJECT

4. Since the June 2, 2026 announcement of Executive Order 14409, "Promoting Advanced Artificial Intelligence Innovation and Security" (the "EO"), Tech Policy Press has sought to provide information to our readership about the frontier artificial intelligence ("AI") "framework" discussed in the EO. Our coverage has consistently noted the gaps in information provided by the White House and other Executive Branch agencies. It has explained that certain AI labs and other companies had privileged information about the "framework" and related efforts that was not available to the public, while attempting to fill in the gaps in the public record through reporting. Tech Policy Press's principal articles on this subject have included:

a. Merve Hickok, “Transparency and Accountability Gaps in Trump’s New AI Executive Order” (June 17, 2026), https://www.techpolicy.press/transparency-and-accountability-gaps-in-trumps-new-ai-executive-order/. This article noted security and oversight risks arising from the limited transparency under the EO, observing that “the EO creates a situation where companies that participate will presumably learn how their own models score, while the public will be kept in the dark. Possibly, even the developers of these models who do not carry security clearance may not know. Even many of the cybersecurity professionals working in major American companies providing critical infrastructure services will not know.” A true and correct copy of this article is attached hereto as Exhibit [A].

b. Michelle De Mooy, “Five Questions the US Government Should Answer About Its Secretive Frontier AI Framework” (Aug. 5, 2026), https://www.techpolicy.press/five-questions-the-us-government-should-answer-about-its-secretive-frontier-ai-framework/. This article placed the opaque EO process in the context of other legislative proposals around AI safety, noting, “Without public disclosure, there is a serious risk that a process built on White House discretion will harden into the permanent structure of frontier AI governance without input from developers, researchers, civil society, or the public.” A true and correct copy of this article is attached hereto as Exhibit [B].

c. Mark MacCarthy, “US Government’s AI Risk Review Should Apply to Open Weight Models” (Aug. 14, 2026), https://www.techpolicy.press/us-governments-ai-risk-review-should-apply-to-open-weight-models/. This piece draws out the implications of the EO regulatory framework for “open-weight” AI models, beginning from the observation that “Little is known about the Trump administration’s secretive new voluntary safety review

program for AI models. Yet one salient detail that is in the public domain is that it apparently excludes open weight models." A true and correct copy of this article is attached hereto as Exhibit [C].

d. Jake Taylor, "AI Systems Are Getting More Powerful. The Ability to Verify Must Keep Pace" (Aug. 25, 2026), https://www.techpolicy.press/ai-systems-are-getting-more-powerful-the-ability-to-verify-must-keep-pace/. The author of this piece was responsible for establishing the U.S. Center for AI Standards and Innovation at the National Institute of Standards and Technology (NIST), and from 2017 to 2020 served as the first assistant director for Quantum Information Science at the White House Office of Science and Technology Policy. He argues that, to close the gap between AI's growing capabilities and the capacity to check these systems, the federal government needs to establish public, standardized testing tools with integrated formal reasoning checks, and explains why building testing standards that are trusted and independent outweighs any asserted risk of disclosure to adversaries. A true and correct copy of this article is attached hereto as Exhibit [D].

5. I personally reported an article on June 13, 2026, describing the interaction between the EO and the export restrictions imposed on two Anthropic PBC AI models: "Anthropic's Mythos Recall and the White House's Missing AI Safety Playbook," https://www.techpolicy.press/anthropics-mythos-recall-and-the-white-houses-missing-ai-safety-playbook/. My article analyzed the "ad hoc and wild swings" in U.S. regulatory policy on frontier AI models, and the consequences of those shifts for international allies and other parties. A true and correct copy of this article is attached hereto as Exhibit [E].

6. This is not incidental coverage. The governance of frontier AI is among the central subjects Tech Policy Press exists to report on, and our audience, much of which is professionally responsible for technology policy, relies on that reporting to do its own work. Articles on this topic have been among our most widely circulated pieces since the publication of the EO.

III. THE UNDISCLOSED FEDERAL REVIEW REGIME

7. Based on public reporting and government statements, I understand the following: (a) an executive order issued June 2, 2026 established a framework under which "covered frontier models" are identified through a classified benchmarking process and reviewed by federal agencies before release; (b) an administration official has stated that a voluntary disclosure and pre-release review framework required by that order was completed by August 1, 2026, but the government has declined to release the framework's text, the identity of participating companies, or its effective date; (c) in June 2026, export-control directives resulted in worldwide suspension of public access to leading frontier models, with restored access to one model initially limited to approximately 100 organizations identified in a non-public annex; and (d) a multi-agency initiative announced July 14, 2026 ("GOLD EAGLE") coordinates AI-related cybersecurity information sharing between the government and industry under undisclosed terms of participation.

IV. HOW THE SECRECY IMPAIRS JOURNALISM ON THIS SUBJECT

8. My ability as a journalist, and Tech Policy Press's ability as a specialized media outlet, to meaningfully report on frontier AI safety regulation has been significantly hamstrung by the limited public information available on the operation of the EO. Tech Policy Press has covered other subjects in which classified government programs were involved, but we have never covered a situation in which a handful of private parties had been given special access to the details of a non-classified, public regulatory structure. Our coverage has struggled to analyze the

implications of this frontier AI regulatory framework in the dark. The paragraphs that follow describe those impairments specifically.

9. The events described above are among the most consequential technology-policy developments of this period, and journalists cannot fully report them. The primary documents that would ordinarily anchor coverage of a new regulatory regime — the framework's text, the criteria for review, the terms of participation, the identity of participants, the basis for selective access — are all withheld. As a result, coverage of this regime, ours included, is necessarily built on officials' characterizations of documents no journalist has seen, on anonymous sourcing, and on inference. That is inadequate to inform our readership that includes AI policy specialists.

10. The consequences are specific. Reporters cannot verify or challenge the government's public statements about the framework against its actual text. We cannot tell readers which companies are inside the regime and which are not, or based on what criteria, even though those facts bear directly on stories about competition, access, and favorable treatment that we would otherwise pursue. We cannot analyze claims that the June directive and its selective restoration reflected published standards, because no standards are published. And when officials or companies decline to answer questions by pointing to the framework's confidentiality, there is no document against which to test that assertion.

11. One line of reporting deserves specific mention. Based on my research and writing on this subject, I am concerned that a secretive framework, affecting fast-growing and extraordinarily wealthy technology companies, creates significant risks of corruption. Tech Policy Press would like to research and report on whether access to the framework and other nonpublic details of the EO's implementation has favored government insiders or benefactors, but we have been unable to pursue that reporting due to the basic lack of transparency around

this program. I do not assert that such favoritism has occurred, only that as matters stand, journalism cannot determine whether it has, which is among the core functions a free press serves.

12. There is also an asymmetry that journalists on this beat confront daily: the framework's contents are known to the participating companies and to the organizations on the non-public access list, while the press and public are excluded. Because sources inside that population face confidentiality expectations, reporting that depends on their willingness to describe documents they are not supposed to describe is a poor and legally fraught substitute for the documents themselves.

13. Based on my research and writing on this subject, I believe the development and application of the EO's secret AI framework is a pivotal moment for AI technology, for the U.S.-based AI labs, for their customers, and for national security. I am deeply concerned that the public and civil society are unable to obtain enough information about the federal government's operation of this framework to meaningfully participate in this important debate.

V. WHY THE RECORDS' VALUE TO JOURNALISM IS TIME-BOUND

14. News value is not durable. The purpose of reporting on a regulatory regime is to inform public and official deliberation while that deliberation can still be affected. Three live windows make that concrete here. First, Congress faces a September 30, 2026 sunset of the statutory information-sharing authorities on which the GOLD EAGLE initiative reportedly depends. Reporting that informs that reauthorization debate must be published before the debate concludes. Second, frontier model releases are reported to be moving through the secret framework's review now. Each release governed by undisclosed criteria is a story that can only be meaningfully reported before or as it happens, not months afterward. Third, the framework itself is still early in its operative phase, when public scrutiny can still inform how it is

implemented. We regularly monitor for any public disclosure or possible evidence of the framework's effects on the behavior of frontier model developers.

15. Tech Policy Press intends to report on the requested records promptly upon their disclosure, including their bearing on the reauthorization debate and on the pre-release review of pending frontier models. In my professional judgment as an editor, records of this kind produced on ordinary FOIA timelines (i.e., many months from now) would arrive after the deliberations they concern have concluded and after the regime they describe has hardened or changed. Their value to journalism, and to the audiences journalism serves, would be largely undercut. For news reporting and analysis, delayed disclosure of time-sensitive government records does not provide a lesser remedy but rather undermines the story's purpose.

16. No party to this litigation compensated Tech Policy Press or me for this declaration, and its contents are my own.

I declare under penalty of perjury that the foregoing is true and correct.

Executed on August 31, 2026, at Brooklyn, New York

Justin Hendrix
CEO and Editor, Tech Policy Press