June 12, 2023

Via Federal eRulemaking Portal

Ms. Stephanie Weiner
Acting Chief Counsel
National Telecommunications and Information Administration
U.S. Department of Commerce
1401 Constitution Avenue NW
Washington, DC 20230

RE: Comment on Artificial Intelligence ("AI") system accountability measures and policies—Docket Number NTIA–2023–0005, 88 FR 22433 (July 12, 2022)

Dear Ms. Weiner,

The undersigned Attorneys General appreciate the opportunity to respond to the National Telecommunications and Information Administration's (NTIA) request for comment on Artificial Intelligence (AI) policies (the RFC). Our comments are informed by State Attorneys General collective and extensive experience enforcing data privacy and consumer protection laws.

I.   Introduction

The NTIA has asked numerous questions about the development of appropriate standards and oversight mechanisms necessary to prevent harms associated with the use of AI. We commend the nature of this inquiry and its commitment to a rigorous and data-driven approach to evaluating the path forward with respect to promoting and protecting trustworthy AI systems. As the RFC recognizes, developing trustworthy AI technologies is a tall order, ensuring, among other things, that AI systems are valid and reliable, safe, secure, and resilient, accountable and transparent, explainable and interpretable, privacy-enhanced, and fair. As the NTIA works to advance this goal, State Attorneys General stand ready to work with you on a range of fronts.

As with other emerging technologies, a critical challenge in this area is to encourage and oversee the proper development of dynamic and trustworthy tools without hampering innovation. This means, for example, that a prescriptive regulatory regime may not be best suited to this challenge. By contrast, commitments to robust transparency, reliable testing and assessment requirements, and after-the-fact enforcement is a very promising approach. We also recommend, as used in the data privacy arena, a risk-based approach, recognizing that certain use cases (say, routes for package delivery) are less concerning than others (say, health care delivery options), though we recognize that a more nuanced evaluation of risk may be required in this context.

The nature of a risk-based approach in regulation is that governments must evaluate and categorize what AI systems could impact, such as collective or individual physical or psychological safety, civil and human rights, or equal access to goods, services, and opportunities.[1] Moreover, levels of risk can also be calibrated by what categories of data are used by the AI system (e.g., sensitive data[2] such as medical information, biometric data, or personal information about children) and what are the ultimate outputs (e.g., "deepfakes"[3] or manipulation). Finally, AI systems could pose higher risk when automated decision-making occurs with no or very limited human agency and those decisions directly impact individuals' legal or financial situation.

Included here are specific recommendations for a governance framework that leverages the public and private sectors and supports the responsible development, use, and deployment of AI systems. These recommendations ensure such systems can develop in a trusted, fair, and technologically dynamic environment.

II. Independent Standards for Transparency, Testing, Assessments, and Audits[4]

The foundation of any effective AI governance framework is appropriate transparency. For example, consumers must be told when they are interacting with an AI rather than a human being and whether the risks of using an AI system are negligible or considerable. To the extent organizations employ AI systems to handle critical functions, the best practice to use is to commit to ongoing cycles of testing, assessment, and external audits. Without such testing, the risk is considerable that they will fail to protect against unintentional and inadvertent harms. Some initial deployments have made clear that AI systems are open to manipulation.[5] Consequently, as part of its work to develop trusted AI systems, NTIA or the National Institute of Standards and Technology (NIST), or both, should work to spearhead consistent criteria and technical standards for testing, assessments, and audit policies and practices of emerging AI systems. Indeed, this testing,

---

[1] NIST's "harm to people" category aligns with this definition, creating three subcategories, including: (a) harm to individual civil liberties, rights, physical or psychological safety, or economic opportunity; (b) harm to a group such as discrimination against a population sub-group; and (c) harm to society, such as harm to democratic participation or educational access. NIST Trustworthy and Responsible AI, National Institute of Standards and Technology, Gaithersburg, MD, [online], https://doi.org/10.6028/NIST.AI.100-1, (Accessed May 22, 2023) (NIST AI RMF 1.0) at 5.

[2] *See*, *e.g.*, Colo. Rev. Stat. § 6-1-1303(24); Va. Code Ann. § 59.1-575; Pub. L. 22-15, 6 § 1(12), Gen. Assemb., Reg. Sess. (Conn. 2022); Cal. Civ. Code § 1798.140(ae).

[3] The term "deepfake" refers to a form of synthetic media, usually an image or video that is altered or manipulated to replace one person's voice or likeness convincingly with that of another. Wikipedia contributors, *Deepfake*, Wikipedia, The Free Encyclopedia (last visited May 18, 2023) https://en.wikipedia.org/w/index.php?title=Deepfake&oldid=1155613928.

[4] This section is responsive to questions 1, 2, 11, 14, 16(b) and 23 of the RFC. 88 Fed. Reg. at 22,435, 22,436, and 22,439-40.

[5] George Petropoulos, *The Dark Side of Artificial Intelligence: Manipulation of Human Behavior*, Brugel Blog (Feb. 2, 2022) https://www.bruegel.org/blog-post/dark-side-artificial-intelligence-manipulation-human-behaviour (last visited June 2, 2023).

rather than stifling development, should be an integral part of the process that leads to further innovation.

A. Transparency

For those organizations engaging in high-risk activities, they should be encouraged or, where appropriate mechanisms (such as procurement requirements) can allow, mandated—to publish public-facing policies that describe what decisions are powered by AI, what human involvement there is in validating those decisions, and what process individuals can use to appeal those decisions. As noted below, key transparency measures should also include information about personal data used in any such decisions and a method for individuals to access and correct any personal information used by the AI in the decisions.

As Justice Brandeis famously said, "sunlight is among the best of disinfectants."[6] By requiring appropriate disclosure of key elements of high-risk AI systems, individuals can be empowered to decide what systems are fair[7] and adhere to critical due process norms.[8] Ultimately, consumers would be well served by the ability of trusted intermediaries to rate and provide guidance on what levels of risk individuals face by the use of different systems. By enabling digestible information to be relied on by regulators and the public, the adoption of AI can increase, and the rights of consumers will be protected.

B. Independent Standards

The NTIA should investigate and establish baseline transparency, testing, assessment, and audit standards. One role of such standards would be to identify the relevant risks associated with AI systems and then work to create a framework for assessing AI systems inputs and outputs relative to the identified risks. Through a subsequent process, the NTIA or NIST, or both, might consider establishing a system for certifying trusted auditors and, in the alternative, a system for establishing and overseeing those entities who can certify trusted auditors.

The development of agile and dynamic public and civic initiatives that build trust and spur trusted technological changes bear consideration by the NTIA as it moves this effort forward. The Energy Star program, for example, was created by the federal government in 1992, adopted by industry, and subsequently codified into statute by Congress.[9] As an

---

[6] Louis Brandeis, *Other People's Money* 92, (Fredrick A. Stokes Co., 1914).

[7] Francois Candelon, Theodoros Evgenious, and David Martens, *AI Can Be Both Accurate and Transparent*, Harv. Bus. Rev. (May 12, 2023), https://hbr.org/2023/05/ai-can-be-both-accurate-and-transparent (last visited May 23, 2023).

[8] This also enables individuals to decide what systems not to use, what systems might be suspect and violative of other regulatory requirements, and what systems are too risky to use. See Margot E. Kaminski & Jennifer M. Urban, *The Right to Contest AI*, 121 Colum. L. Rev. 7, 1957, 1979 (Nov. 2021).

[9] Energy Star, *How ENERGY STAR works: Our History*, https://www.energystar.gov/about/how_energy_star_works/history#:~:text=The%20U.S.%20Environmental%20Protection%20Agency,go%20hand%2Din%2Dhand (last visited June 1, 2023).

example of a private sector program, the Leadership in Energy and Environmental Design (LEED) standard has spurred the move towards "green buildings."[10]

The NTIA has an opportunity to lead in the emergence of trusted AI systems by developing a respected governance regime. To do so, the NTIA could follow up this RFC by convening industry, independent standards-setting bodies, governmental leaders, and other AI stakeholders to develop the architecture for such a program. The ultimate goal of this program would be to foster similar trust in AI systems through transparent and verifiable policies and practices driven by appropriate standards including a code of ethics. For companies who adopted and committed to such practices, moreover, the Federal Trade Commission and State Attorneys General would possess the ability to enforce the commitments companies made using their consumer protection authority (which sanctions deceptive trade practices). This legal architecture would parallel that used to develop oversight of data privacy policies adopted by companies as well as that used in the cybersecurity space, where some entities publicize their compliance with NIST, SOC II, or ISO standards to indicate their compliance with best practices and applicable law.

The development of appropriate standards for trusted AI should flow from a multi-stakeholder process that is transparent, inclusive, and accessible, and incorporates a code of ethics.[11]. To architect such a model, the NTIA should investigate what ingredients guided past efforts that built trust and spurred technological advancement in different fields (such as energy-efficient products or green buildings).[12] Ultimately, like Energy Star and LEED, it will be critical to establish a system for certifying compliance with the relevant standards. And like the initial steps forward in data privacy and data security, the FTC and State AGs promise to play an important role in ensuring that organizations adhere to their public commitments.

C. Internal Testing and Impact Assessments

In addition to transparent disclosure of how AI systems work and the development of standards for trusted AI, entities must engage in periodic testing and assessment of AI systems that pose a notable risk to the legal rights of consumers or citizens. This means, for example, that where safety, financial consequences, or illegal discrimination are reasonably possible on account of AI systems, those implementing such systems should adopt regular testing and auditing that evaluates such possible risks. In the case of privacy protection, for example, such approaches, including those known as "privacy by

---

[10] *See LEED Rating System*, https://www.usgbc.org/leed (last visited May 23, 2023).

[11] *See, e.g.,* V.S.A. § 5022 (b)(1) and Final Report, Artificial Intelligence Task Force, at 17-19, Jan. 15, 2020, *available at*: https://outside.vermont.gov/agency/ACCD/ACCD_Web_Docs/ED/MajorInitiatves/ArtificialIntelligenceTaskForce/FinalReport.pdf.

[12] For one discussion of this effort, see Philip J. Weiser, *Entrepreneurial Administration*, 97 B.U. L. Rev. 2011 (2017).

design,"[13] have improved the overall design and development processes of data management.

It is important for impact assessments to be ongoing and begin as early as the design phase of a new AI system and even before an AI system is put into use. Such assessments can be useful in identifying potential harms and allowing developers to make design changes to avoid these harms and create a baseline level of understanding. Consequently, further periodic assessments and testing conducted once the system is operational can reveal unexpected results or unintended consequences that emerge over time. Such learnings can enable developers and users of AI to remediate and mitigate unintended harms on an ongoing basis.

The very nature of the development and implementation of AI will both create and mitigate risks as to how operations were managed previously. It is critical that the designers and operators of AI systems develop a learning mindset, building in regular evaluation and testing to understand and manage the relevant risks at every phase of development. As they do so, they should evaluate to what extent risks can be mitigated or emerge at different phases of development and operation.

For an example of how an assessment process can operate, consider the Colorado Privacy Act (CPA) rules' detailed provisions for the contents of Data Protection Assessments.[14] At its core, the CPA requires that Data Protection Assessments contain a "genuine, thoughtful analysis."[15] By design, these assessments focus on substance over form, encouraging a thoughtful assessment rooted in the particular use case and risk of harm.

Significantly, an assessment or audit need not be overly complex. An AI system assessment might, for example:

1. Identify and describe the specific risks to safety or individual or collective civil and human rights;
2. Identify risks associated with the data used by the system (e.g., biometric, financial, PII, or large-scale generalized data);
3. Document measures taken to avoid or offset those risks;
4. Document "grounded" tests or otherwise independent testing of the AI system to demonstrate efficacy without unintended bias, errors, or false outcomes;
5. Contemplate the benefits of the AI system; and
6. Demonstrate that the benefits of the system outweigh the risks offset by safeguards in place.

---

[13] *See* Ann Cavoukian, Ph.D, *Privacy by Design: The 7 Foundational* Principles, Information and Privacy Commissioner of Ontario, https://www.ipc.on.ca/wp-content/uploads/resources/7foundationalprinciples.pdf (2011) (Establishes a principled approach to ensure privacy is considered at all phases of design processes and product lifecycles). See also Regulation 2016/679, art. 25, 2016 O.J. (L 119/1) (EU) (General Data Protection Regulation).
[14] 4 Colo. Code Regs. § 904-3-8.
[15] 4 Colo. Code Regs. § 904-3-8.02(A).

To incentivize true reflection and a learning mindset, we do not believe that these assessments must be subject to public scrutiny or be made public as a default. As is required under the CPA, however, enforcement authorities should have the right to request these assessments and audits at least annually to ensure that such assessments take place with the appropriate vigilance and are not an empty promise. To the extent that assessments or audits are required (say, as part of procurement contracts or regulatory requirements), they should be accompanied by appropriate record-keeping requirements.

D. External, Third-Party Audits

Finally, entities that use or develop high-risk AI should be required to engage in periodic external, third-party audits to ensure that any AI systems in use by the entity comply with a common set of criteria set by independent standards. Ideally, these external audits would build on the internal assessment systems noted above, providing external oversight that validates internal testing and assessment as well as ensuring that entities operate with a mindset of risk assessment and continuous improvement. To optimize the effectiveness of such audits, the NTIA, NIST, or another trusted standard-setting body could develop appropriate protocols for such audits.

III. Legislation and Enforcement of AI Governance Standards[16]

A. Federal Legislation

The development of principle-based governmental oversight of AI will be imperative in the years ahead. There are a number of potential paths that such oversight could take, including potential federal legislation.[17] Whatever form such oversight takes, it is important that it recognize the potential risks that AI systems could pose and ensure that they operate in a responsible and trustworthy fashion. This oversight must also recognize the dynamic nature of AI and avoid the mismatch of overly prescriptive requirements on these technologies. Indeed, proper regulation should recognize and incorporate proper incentives to foster innovation and better, safer outcomes.

A focus on high-risk AI systems aligns with existing AI regulatory frameworks, such as the EC AI Act, which prohibits the use of AI in certain high-risk contexts[18] and imposes heightened obligations in other such systems.[19] U.S. privacy law has created similar

---

[16] This section is responsive to questions 17, 25, 30, 32 and 33 of the RFC. 88 Fed. Reg. at 22,440.

[17] One approach would be to empower an agency with specific expertise to oversee Internet platform technologies. See, e.g., Letter from Attorneys General Josh Shapiro and Philip J. Weiser to the Honorable Marsha Blackburn and the Honorable Richard Blumenthal (October 6, 2021) available at https://coag.gov/app/uploads/2021/10/Internet-Regulation-Letter-to-US-Senate-10.0.21-Final.pdf.

[18] *Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts*, Title II, Article 5 at 43, COM (2021) 206 final (Apr. 21, 2021).

[19] *Id.*at 45 (Stating that an "AI system shall be considered high-risk where both of the following conditions are fulfilled: (a) the AI system is intended to be used as a safety component of a product, or is itself a

guardrails around high-risk use cases. Laws in Colorado, Connecticut, Montana, and Virginia, for example, set heightened requirements when data is processed to support decisions that result in legal or other significant effects for an individual, such as automated decisions that impact an individual's access to financial, educational, housing, or employment opportunities.[20] Laws governing facial recognition technologies have also created heightened requirements in high-risk contexts, requiring meaningful human review when government entities use facial recognition technologies in ways that may impact individual liberties and rights.[21]

Building a framework emphasizing high-risk use cases instead of all development and use of AI is more likely to protect against the most egregious harms while leaving space for ongoing innovation and appropriate self-regulation for lower-risk uses. In addition, any legislation should, at a minimum, require human review of AI-driven decisions in the high-risk context and impose robust transparency requirements on developers, operators, and users of AI systems that have a notable potential impact on consumers. Individuals should also be granted the right to appeal any decisions fostered by AI and the right to correct any personal information that is fed into an AI system.

B. Privacy Legislation and AI

While the lack of federal privacy legislation is not necessarily a barrier to creating effective AI accountability, any AI governance standards and regulations should consider, and be interoperable with, existing privacy regimes aimed at protecting individual privacy rights. Specifically, any AI governance standards should align with existing state and federal data security requirements to ensure that personally identifiable information used to inform AI systems is adequately secured.

In addition, any legislation aimed at AI should contemplate the privacy rights impacted by the data required to power effective AI, including appropriate consent and transparency requirements both for the initial collection and use of sensitive data used by AI systems, as well as the repurposing of personal data collected by companies for other reasons to help inform AI systems. Regulators of AI should be mindful of existing privacy laws and create standards and legislation that build off existing regimes to avoid any conflicts.

---

product, covered by the Union harmonisation legislation listed in Annex II; (b) the product whose safety component is the AI system, or the AI system itself as a product, is required to undergo a third-party conformity assessment with a view to the placing on the market or putting into service of that product pursuant to the Union harmonisation legislation listed in Annex II."). Annex II includes a long list of identified high-risk AI systems.

[20] *See* Colo. Rev. Stat. § 6-1-1303(10) (2023) (Colorado Privacy Act), Conn. Gen. Stat. § 42-515 *et seq.* (Connecticut Data Privacy Act), Legislature Regular Session (Mont. 2023) S.B. 384 (Montana Consumer Data Privacy Act), § Va. Code Ann. § 59.1-575 (2023) (Virginia Consumer Data Protection Act).

[21] *See* Colo. Rev. Stat. § 24-18-301(3) (2022) (Colorado Law Restricting the Use of Facial Recognition Services by State and Local Government Agencies), Wash. Rev. Code Ann. § 43.386.030 (West) (Washington Law Concerning the Use of Facial Recognition); Davis, Ca., Mun. Code art. 26.07; Nashville, Tenn., Ordinance No. BL2017-646 (June 7, 2017).

C. Enforcement

State Attorneys General should have concurrent enforcement authority in any Federal regulatory regime governing AI. Significantly, State AG authority can enable more effective enforcement to redress possible harms. Consumers already turn to state Attorneys General offices to raise concerns and complaints, positioning our offices as trusted intermediaries that can elevate concerns and take action on smaller cases.

The current reality in the data privacy arena is one where states are active and engaged in protecting consumers. The Colorado Privacy Act (CPA), [22] Connecticut Data Privacy Act (CTDPA),[23] California Privacy Rights Act (CPRA),[24] Tennessee Information Privacy Act (TIPA)[25], and Virginia Consumer Data Protection Act (VCDPA)[26] also specifically regulate AI systems through statutes and rules which govern profiling and automated decision-making. Under the CPA, for example, businesses must enable consumers to opt out of the processing of personal data for any automated decisions that produce legal or other significant effects for a consumer, such as automated decisions that impact financial, educational, housing, or employment opportunities.[27] Under current law, businesses are required to conduct data protection assessments before carrying out automated data processing, which carries a risk of unfair treatment or financial, physical, or other substantial harm to consumers.[28] Allowing for concurrent state Attorney General enforcement of any federal AI regulation would allow for continued enforcement of these and other important protections regarding AI. Concurrent enforcement will also expand the resources available to review risk assessments like those required by the CPA and the CTDPA.

Finally, whatever accountability mechanism is developed, it should ensure avenues for legal redress against all participants within the chain of use and development of an AI System that causes legally cognizable harms.

* * *

Thank you again for the opportunity to provide comments in response to these timely questions. We look forward to reviewing any standards or recommendations the NTIA puts forward based on the responses it receives to the RFC and working with you on this important initiative.

---

[22] Colo. Rev. Stat. §§ 6-1-1302(20), 6-1-1306(1)(a), 6-1-1309; 4 Colo. Code Regs. §§ 904-3-9.
[23] Conn. Gen. Stat. § 42-515 *et seq.*
[24] Cal. Civ. Code §§ 1798.100-199.
[25] Tenn. Code Ann. § 47-18-3201.
[26] Va. Code Ann. § 59.1-575 (2023).
[27] Colo. Rev. Stat. § 6-1-1306(1)(a)(I)(C); 4 Colo. Code Regs. § 904-3-9.04; *see also* Conn. Gen. Stat. § 42-518 (a)(5)(C).
[28] Colo. Rev. Stat. §6-1-1309; 4 Colo. Code Regs. § 904-3-9.06; *see also* Conn. Gen. Stat. § 42-522.

The four co-sponsors of this letter, the attorneys general of Colorado, Connecticut, Tennessee, and Virginia are joined by the undersigned attorneys general across the U.S. states and its territories.

Sincerely,

Phil Weiser

Colorado Attorney General
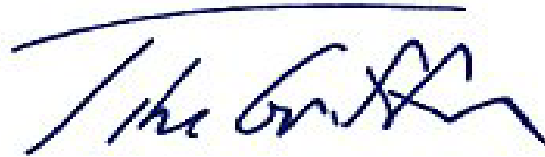
William Tong

Connecticut Attorney General

Jonathan Skrmetti

Tennessee Attorney General

Jason S. Miyares

Virginia Attorney General

Kris Mayes

Arizona Attorney General

Tim Griffin

Arkansas Attorney General

Rob Bonta

California Attorney General

Kathleen Jennings

Delaware Attorney General

Brian Schwalb

District of Columbia Attorney General

Kwame Raoul

Illinois Attorney General

Aaron M. Frey

Maine Attorney General

Keith Ellison

Minnesota Attorney General

Aaron D. Ford

Nevada Attorney General

Matthew J. Platkin

New Jersey Attorney General

Letitia James

New York Attorney General

Josh Stein

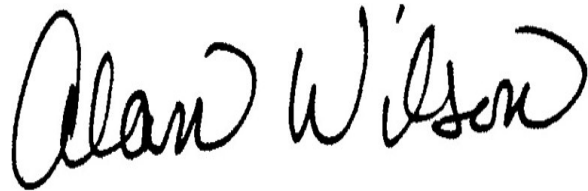North Carolina Attorney General

Dave Yost

Ohio Attorney General

Gentner Drummond

Oklahoma Attorney General

Michelle Henry
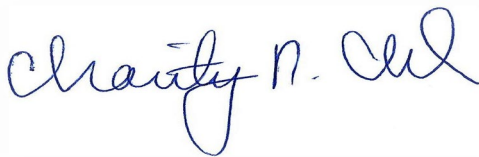
Pennsylvania Attorney General

Alan Wilson

South Carolina Attorney General

Marty Jackley

South Dakota Attorney General

Ariel M. Smith

U.S. Virgin Islands Acting Attorney General

Charity Clark

Vermont Attorney General