credo ai

March 15, 2025

To: Faisal D'Souza
Office of Science and Technology Policy
Executive Office of the President
2415 Eisenhower Avenue
Alexandria, VA 22314

RE: OSTP Request for Information on the Development of an Artificial Intelligence (AI) Action Plan Federal Register #: 2023-11346

This document is approved for public dissemination. The document contains no business-proprietary or confidential information. Document contents may be reused by the government in developing the AI Action Plan and associated documents without attribution.

About Credo AI

Founded in 2020, Credo AI's mission is to empower enterprises to responsibly build, adopt, procure and use AI at scale. Credo AI's pioneering software AI governance platform helps enterprises from the Global 2000s to small- and medium- sized enterprises measure, monitor and manage AI risks to enable faster trusted AI adoption.

Credo AI's work aligns with the Administration's goal of promoting innovation and seizing the incredible opportunities presented by AI to secure and strengthen America's leadership in Artificial Intelligence. We believe that building trust in this technology will encourage the broader adoption of AI across all sectors and markets, thus accelerating AI innovation and our leadership. Our Credo AI platform operationalizes industry frameworks, standards from international organizations, and U.S. guidance (such as the NIST AI Risk Management Framework) so that organizations adopt and build AI with speed and at scale in alignment with their goals.

Executive Summary

Credo AI welcomes the opportunity to contribute to the formulation of the AI Action Plan pursuant to Executive Order 14179, "Removing Barriers to American Leadership in Artificial Intelligence." In our work to enable AI adoption with organizations, we have seen a clear need for certainty in the ecosystem in order to enable organizations to act faster and capitalize on the gains of AI at scale, without suffering brand or reputational damage due to unrecognized or unmitigated AI risks. Our recommendations focus on how to build trust and confidence in AI systems, and create certainty for businesses to make their own informed choices, enabling them to build faster, better, trusted AI and accelerate AI innovation. By ensuring secure and transparent deployment of AI, American enterprises will be better positioned to leverage these capabilities to accelerate innovation. To facilitate this, we encourage the Trump Administration's AI Action Plan to highlight and actively support the following priorities:

1. **Scale Responsible AI usage to increase U.S. Government efficiency.** Enable U.S. agencies to evaluate the security and safety of AI models that are procured or deployed for public services, increase the availability of AI tools and training for government employees, and encourage new entrants and SMEs into the procurement process to make federal agencies more efficient and cost-effective.
2. **Invest in public-private partnerships to develop robust evaluations and benchmarks.** Develop robust model evaluations and benchmarks to ensure the responsible and effective deployment of powerful AI systems. A critical component of this effort is the development of context-specific use focused benchmarks, addressing a current gap in the evaluation of AI models. To achieve this, we urge the government to prioritize the creation and adoption of targeted benchmarks, including benchmarks for national security use of AI.
3. **Ensure transparency of the AI value chain to accelerate adoption and protect national security.** Continue the development and use of transparency measures throughout the entire AI value chain, in order to

both accelerate enterprise adoption and enhance the ability of the U.S. government (USG) to generate and request evidence pertaining to national security concerns, especially related to foreign AI models. Both U.S. consumers and the USG need to trust in AI in order to encourage its widespread adoption and use. Trust can be built with and between American businesses using risk mitigation techniques and disclosure mechanisms.

4. **Sustain American leadership in technical AI standards development.** Continue engaging at the international level to advance global standards for AI development and deployment, to give enterprises the ability to effectively compete in markets around the world. American enterprises have repeatedly advocated for global alignment around AI standards and benchmarks to better enable their ability to compete in various markets. Recognizing the interdependence of the AI value chain, U.S. leadership is needed in international engagement and advocacy for adherence to technical standards can help to align conversations about AI governance and transparency, and promote an open-innovation-oriented approach for enterprises globally.
5. **Prioritize U.S. Leadership in open-source and open-weight models.** To remove barriers to adoption, enterprises and the USG need to be enabled and prepared to make independent decisions about business partners in a globally interconnected AI ecosystem. To do this, both enterprises and the USG need to understand to what extent a model is "open." We encourage the USG to lead discussions that move stakeholders to a consensus understanding of what constitutes a truly open source or open weight AI model, including the governance of these models.

By implementing these measures collectively —scaling responsible AI to increase government efficiency, investing in robust evaluations and benchmarks, ensuring transparency of the AI value chain, sustaining American leadership in technical AI standards development, and prioritizing discussions on open-source and open-weight models —USG and enterprises will be more confident in adopting trusted AI, spurring further innovation while

mitigating key risks and aligning to USG and organizational values. Below, we expand on specific policy recommendations that can help to enhance both the USG and enterprises' ability to build, adopt, procure, and use trusted AI at scale.

1. Scale Responsible AI Usage to Increase Government Efficiency

Many steps have already been taken to modernize the government and increase efficiency using AI, including replacing outdated IT systems with cost-effective AI solutions for various internal services such as HR, payment processing, and claims management. To build on this success, the USG should further prioritize AI-driven modernization, supplementing manual processes with trustworthy, secure AI systems. These steps would streamline government operations, enhance security, and drive efficiency across federal agencies.

As part of these efforts, the we recommend that the USG:

- **Create a procurement environment that balances innovation and responsible governance.** Integrate rigorous security and reliability standards into Federal procurement and acquisition of AI vendors, including foundation guidelines for what Federal agencies should evaluate when procuring AI solutions to speed up the procurement and implementation cycle.
- **Invest in AI governance talent and AI literacy within the Federal Government** to ensure procurement moves fast while preventing risks that can delay processes or delivery of essential government services.
- **Encourage new entrants into the federal procurement process, including SMEs and cutting-edge AI start-ups** that are leaders in AI innovation, with new age AI literate teams that can be cost efficient solutions to ensure U.S. competitiveness and leadership.

2. Invest in Public-Private Partnerships to Develop Robust Evaluations and Benchmarks

We recommend the implementation of robust evaluations and benchmarks to ensure the responsible and effective deployment of powerful AI systems. A critical component of this effort is the development of context-specific benchmarks, addressing a current gap in the evaluation of AI models.

To achieve this, we urge the government to prioritize the creation and adoption of targeted use case focused benchmarks, similar to Credo AI's Model Trust Scores, in collaboration with organizations such as MLCommons, Scale AI, and other organizations that have done a significant amount of research in this space. This initiative should build upon the foundational work of NIST's AI Risk Management Framework (AI RMF), and advance the use case-specific evaluations and benchmarking methodologies. By ensuring that evaluations account for use case specificity, the U.S. government can enhance trust and accountability in the deployment of AI systems.

Additionally, incorporating red teaming and automated assessments is essential to strengthening AI governance. All evaluations must be embedded within a comprehensive governance framework that facilitates holistic risk management while aligning with the government's internal values and regulatory priorities.

To facilitate this, we recommend that the USG:

- **Continue investing in the National Institute of Standards and Technology (NIST), and utilize the existing work developed and led by NIST, such as the NIST AI Risk Management Framework (RMF), along with the developmental work conducted by the U.S. AI Safety Institute (AISI),** to create incentives for AI providers to provide baseline documentation and results of evaluations on how models are developed (e.g., model and dataset cards), and, facilitate industry and public sector adoption of AI.

- **Support standardized AI validation protocols for adversarial testing and red-teaming to evaluate AI robustness, security, and non-negotiable requirements for enterprise adoption.** Enterprises are looking for certainty in the models they are integrating into their businesses to ensure that their use of AI systems aligns with their goals. Non-negotiable baseline requirements that industry is looking at when making decisions about incorporating AI into their business include security, reliability, and compliance. Having information to evaluate and make informed decisions about what models best meet these baseline requirements and fit their business goals is critical to accelerate AI adoption and innovation.
- **Build upon the work conducted by organizations like Credo AI, MLCommons, Scale AI and others to develop targeted, context-specific benchmarks and evaluations for AI models.** As AI capabilities continue to accelerate and available model options increase, use case focused benchmarks become critical in understanding tradeoffs - how to make the AI model selections fit for your business purpose. Benchmarks provide a reference point and comparison for enterprises, and increase the utility of evaluations that can serve as a starting point for marketplace-wide standards. We urge the government to prioritize the creation and adoption of targeted use case focused benchmarks, similar to Credo AI's Model Trust Scores, in partnership with leading academia, public and private industry organizations.

3. Ensure transparency of the AI value chain to accelerate adoption and protect national security

As AI capabilities continue to increase, it becomes more critical to have clear AI transparency mechanisms to lead. Transparency along the AI value chain and stack is critical to maintaining the integrity of AI systems, especially in relation to the use of adversarial and foreign models. If done right, transparency along the AI stack contributes to innovation and competition by broadening access for businesses. Enterprises across the value chain benefit from having access to information about the models, vendors that are

powering an end AI application, and how foundational models are trained, in order to determine whether they can safely build on these modes.

- **To create trust, companies should conduct internal risk assessments, strengthen their internal documentation, and fortify their record keeping of AI system functionality and development processes.** Enterprises can and should be identifying and mapping all AI systems in use, flagging high-risk use cases, and enabling a plan of action to modify or potentially suspend prohibited use cases or prohibited adversarial and foreign models. Doing this will help to ensure that American businesses are able to be transparent about what models, datasets, and vendors they interact with in the "AI stack," and account for any variations in security measures.
- **Disclosure along the AI value chain can also be utilized to encourage trust as part of transparency.** Disclosure reporting and mechanisms like system cards, model cards and impact assessments build trust with consumers, along with dataset transparency used by open source and proprietary models.

The proliferation of powerful AI necessitates baseline transparency and robust evaluations for businesses to effectively identify and address emergent threats, including adversarial vulnerabilities, algorithmic bias, privacy breaches, and socio-economic disruptions. Such disruptions can be costly to U.S. businesses integrating powerful AI models into downstream applications.

4. Prioritize U.S. Leadership in Discussions on Open-Source and Open-Weight Models

With more certainty in this field, enterprises would be better enabled to make their own decisions on business partners within the AI ecosystem. Focusing these discussions on what information is necessary to share downstream will allow enterprises throughout the value chain to determine whether or not it is safe and secure for them to be building upon and working with various model providers, and in which markets.

In order to advance these discussions, we recommend that the USG:

- **Encourage public-private partnerships with industry to address concerns related to open-source and open-weight models, and aim to align on an understanding of what it means for a large language model to be "open."** We encourage the U.S. government to lead discussions that move stakeholders to a consensus understanding of what constitutes a truly open source or open weight AI model, including the degree to which original training data sets for models are made available.
- **Ensure that discussions around AI openness include considerations of data provenance, model transparency, and security protocols** to help mitigate concerns about unintended vulnerabilities or regulatory misalignment. A well-defined framework for openness will not only enhance trust among stakeholders but also support innovation by enabling enterprises to make informed, strategic decisions about integrating AI within their operations. The U.S. government's role in facilitating these conversations will be critical to ensuring alignment between industry players and policymakers, ultimately driving the responsible and secure adoption of AI technologies.

After defining open-source, it is critical to have the right structures in place for AI used in government applications to ensure appropriate oversight. The ability of enterprises to move faster with AI will depend on enabling businesses to independently make their own choices about their business partners (and models).

5. American Leadership in Technical AI Standards Development

Because the technological conversation is not bounded by nation state borders, it is more important than ever that the U.S. government continues to engage in conversations about the future of AI development on the global level. This will help to ensure approaches are interoperable, and that countries and organizations are best leveraging their strengths to advance safe and

secure AI adoption aligned with state-of-the-art technological developments.

Public and private sector stakeholders should also keep in mind the interdependence of the AI value chain and the necessity of international engagement. Because AI R&D, supply chains, and data sharing transcend national borders, consistent engagement on AI governance and transparency is essential. Aligning standards and risk mitigation approaches with global partners can help countries and organizations leverage their respective strengths and keep advancing the state-of-the-art in meaningful ways.

Regulatory structures can facilitate an innovation-conducive environment. The National Institute of Standards and Technology (NIST) existing work such as the AI Risk Management Framework has enabled the U.S. to lead the risk management conversation globally. Further collaboration and investment in technical standardization and harmonization can catalyze innovation while concurrently addressing latent risks.

The U.S. is a unique position to advance this work with the following actions:

- **Enhancing AI safety and performance benchmarks** by building on NIST's efforts and defining standardized evaluation metrics encompassing robustness, fairness, and security especially at use case and AI application level.
- **Continuing cooperation with like-minded partners** to harmonize AI governance approaches, reducing regulatory divergence that may hinder U.S. AI companies with global operations.
- **Investing in sector-specific AI governance frameworks**, especially in high-risk sectors. This would enable enterprises with high-risk applications such as healthcare, finance, and national security to leverage uniform and interoperable frameworks and increase certainty that their AI systems behave in alignment with their goals.
- **Maintaining international strategic alliances** to ensure that USG agencies (State Department, Commerce Department, and others)can

continue to promote US industry-driven, consensus-based standards and frameworks (such as the NIST AI RMF) and promote transparent, innovation focused AI governance standards.

Conclusion

Taken together, these efforts —scaling responsible AI to increase government efficiency, investing in robust evaluations and benchmarks, ensuring transparency of the AI value chain to mitigate use of adversarial and foreign models, sustaining American leadership in technical AI standards development, and prioritizing discussions on open-source and open-weight models — will ensure that the USG's AI Action Plan does not stifle innovation, but rather catalyzes it by clarifying responsibilities and creating a stable environment for AI growth.

Credo AI remains steadfast in its commitment to collaborating with the Trump administration, OSTP and the broader USG to architect AI governance frameworks that simultaneously accelerate AI innovation, reinforce our international AI leadership, and ensure secure trusted AI development and deployment. Governance—when strategically structured—not only enhances trust and accountability, but also serves as a critical accelerant for trusted AI adoption for both USG and enterprise competitiveness. We appreciate the opportunity to contribute to this pivotal initiative, and look forward to our continued engagement with the USG on this critical initiative.

Respectfully,

Navrina Singh
Founder and CEO
Credo AI
www.credo.ai

Evi Fuelle
Global Policy Director
Credo AI

Credo AI
4546 El Camino Real B10 #795
Los Altos, CA 94022