600 14th St. NW, Suite 300
Washington, D.C. 20005

February 28, 2025

Subject: **RFI on the Development of an Artificial Intelligence Action Plan**

To Whom It May Concern:

On behalf of International Business Machines Corporation (IBM), we welcome the opportunity to respond to the National Science Foundation's (NSF) and Office of Science and Technology Policy's (OSTP) request for information (RFI) regarding the "Development of an Artificial Intelligence Action Plan." IBM wholly supports the Trump Administration's recognition "that with the right government policies, the United States can solidify its position as the leader in AI and secure a brighter future for all Americans."

To achieve that end, we recommend the Trump Administration's AI Action Plan highlight and actively support the following four priorities:

- **Accelerate Adoption of American Open Source AI.** Enable open innovation in AI by building on past successes to build shared computing and data resources that can serve as the infrastructure for open innovation ecosystems, as well as incentivizing AI skills in the workforce.

- **Formalize an "AI Diplomacy" Strategy to Lead Globally.** Deputize the Office of the Special Envoy for Critical and Emerging Technologies as the lead agency tasked with developing a comprehensive global engagement strategy for supporting American AI interests, and focus export control restrictions on hardware, not AI model weights. Additionally, task the National Institute of Standards and Technology (NIST) to continue promoting US industry-driven, consensus-based standards and frameworks for AI to advance security and trust and to enable global adoption. Features of such a strategy should also include:

    - Protecting US market access and promoting an open-innovation-oriented approach globally, including by advocating for adherence to technical standards.

    - Pushing back against efforts that would undermine the ability of American companies to rely on flexible and reasonable IP frameworks, such as the fair use doctrine, for AI training and development.

- **Prioritize Risk-Based Approaches to Regulation.** Collaborate with Congress to advance legislation that preempts the emerging patchwork of state legislation on AI. Such legislation should refrain from mandating third-party AI audits and instead focus on light touch requirements to improve transparency and documentation, and address gaps in the application of existing law to high-risk uses of AI, where appropriate.

- **Rapidly Scale AI Usage to Streamline Government Operations.** Scale machine learning, generative AI, and automation that are already embedded in commercially hosted shared services platforms to make government agencies more efficient. As envisioned in OMB Directives issued during the first Trump Administration, prioritize rapid migration of costly, bespoke legacy IT systems to cross government platforms in areas like HR, travel, and payment processing.

IBM commends NSF and OSTP for seeking broad stakeholder input on the development of the administration's new AI policy agenda. Pursuing an AI policy agenda that adopts a permissive regulatory regime will ensure a flourishing American Golden Age of technological innovation and economic prosperity.

As you proceed, we welcome the opportunity to continue engaging with the administration to promote a responsible approach to AI that protects consumers without stifling innovation.

Respectfully,

Christina Montgomery
Chief Privacy and Responsible Technology Officer
IBM Corporation

IBM Response to *Request for Information on the Development of an Artificial Intelligence Action Plan*

IBM commends the Trump Administration for taking the pivotal first step in developing a comprehensive AI Action Plan to help ensure continued American leadership in the development and deployment of AI models and systems.

Accelerate Adoption of American Open Source AI

Openness has always been a force for improving safety and security. For decades, increased transparency has enabled a broad and diverse stakeholder community to identify and fix vulnerabilities, increase performance and resilience, and raise the bar for security overall. Further, open model weights will be key to enabling new technical breakthroughs that increase safety. For example, researchers developed a "self-destructing models" technique that reduces the risk of bad actors exploiting dual-use models.[1] This work was only possible because they had access to open model weights for cutting-edge models.

Open models encourage greater involvement and scrutiny from a larger community of stakeholders, increasing the likelihood that bias and vulnerabilities are identified and patched. Ultimately, when a broad community is represented in AI development, safety is prioritized. These benefits are not confined to the commercial sector, but extend to the national security and intelligence community as well.

The recent National Telecommunications and Information Administration (NTIA) report on open model weights confirms these benefits and affirms the value and safety of open AI technologies for innovation. It concluded that risk is not inherent to open or closed AI models and that "government should not restrict the availability of open model weights."[2] The risks that could be influenced or exacerbated by open access to model weights are not well defined or backed by robust, transparent scientific evidence. Any conclusions about how to best manage these risks are necessarily premature. Non-governmental actors have similarly confirmed these findings. For example, RAND published a report noting that leading AI models did not provide any advantages to bad actors in the domain of biological risk that could not be gained from simple Google searchers.[3]

While there is a lack of evidence around risk of open access to model weights, the broad economic and social benefits of openness are overwhelmingly evident. Those benefits, in turn, create an economic powerbase upon which national security investments rely. A strong economy is a

[1] Peter Henderson, et. al., "Self-Destructing Models: Increasing the Costs of Harmful Dual Uses of Foundation Models," *arxiv*, 27 Nov. 2022, available at https://arxiv.org/abs/2211.14946.
[2] *Dual-Use Foundation Models With Widely Available Model Weights Report*, National Telecommunications and Information Administration, 30 July 2024, available at https://www.ntia.gov/programs-and-initiatives/artificial-intelligence/open-model-weights-report.
[3] "Red-Teaming the Risks of Using AI in Biological Attacks," RAND, 25 Mar. 2024, available at https://www.rand.org/pubs/articles/2024/red-teaming-the-risks-of-using-ai-in-biological-attacks.html.

necessary precondition for a strong national security, and advancements in AI made possible by open source developments are a critical cornerstone of modern American economic competitiveness.

Openness also lowers the barriers to market entry. Companies should compete based on how well they can tailor and deploy AI in valuable ways – not on how efficiently they can hoard. For example, by making many of the technical resources necessary to develop and deploy AI more readily available (including model weights), open ecosystems enable small and large firms, start-ups, and research institutions to develop new and competitive products and services without steep, and potentially prohibitive, upfront costs. And there is tremendous demand from businesses for these kinds of resources.

Sacrificing an open ecosystem for AI in the name of safety would give only a select few firms the opportunity to hoard the benefits of innovation and fail to develop meaningful and sustainable solutions for potential safety and security risks that may arise. Open models and proprietary models can, and should, coexist within the broader AI marketplace, and businesses and consumers should be able to make their own choice.

We would not be at this AI moment if it were not for the diverse scientific and technical community that has openly contributed for decades to the fundamental advances that we benefit from today. Open access to AI means more stakeholders can identify opportunities to improve AI technologies and more easily pursue novel and valuable applications for AI. To ensure innovation, security, and choice, policymakers should avoid putting a thumb on the scale against open.

Recommendations

- Domestically, the Trump Administration should enable open innovation in AI by building on past successes and building shared computing and data resources that can serve as the infrastructure for open innovation ecosystems. This could involve, for example, developing a National Compute Strategy, funding the National AI Research Resource, and better coordinating public-private initiatives to make it easier to access the computing power necessary to develop and evaluate AI.

Formalize an "AI Diplomacy" Strategy to Lead Globally

National Security requires not only strength in arms, but also the ability to leverage soft power through technological and economic channels. Over the past half decade, technology diplomacy has grown in importance, and is now on par with trade and military cooperation in international negotiations, forming the third leg in America's diplomatic "triad" toolkit. A robust technology sector confers this advantage on the exercise of American soft power abroad, and open source AI plays a critical role in accelerating the American technology sector's ability to remain on the cutting edge of innovation relative to foreign competitors.

AI policy is a new part of America's diplomatic "triad." The Trump Administration should leverage America's competitive advantage in this space and ensure the proliferation of policies that promote an open AI future. As Sen. Todd Young recommended in a recent *National Interest* essay, the administration

> "should consider ways to present a better option to a world demanding technology solutions by harnessing the power and capacity of America's technology sector. This could be a formalized program for exporting U.S. technology solutions in computing, AI, or biotechnology with certain incentives and expectations attached, similar to how we approach the export of defense articles through the Foreign Military Sales (FMS) process."[4]

Additionally, Sen. Young suggests opportunities for considering how to "extend the recent successes in semiconductor diplomacy to other emerging technologies in order to promote openness, democracy, and the rule of law."[5] These recommendations all point towards the need for a concerted "AI Diplomacy" strategy – one that anchors itself on the need to promote open source AI diffusion abroad.

Open source AI diffusion is critical for American success in the global race for technological supremacy, as the nation that maintains leadership in the development or maintenance of the open source AI ecosystem will have a considerable influence on the global stage. The more developers and researchers that build the next generation of AI on American open source models, the more the geopolitical concentration of technological power will reside within America's sphere of influence. This influence will be critical for ensuring the continued technological dominance of the democratic West, while undermining the growing threat of techno-authoritarian values.

If American developers are unduly burdened by government-imposed constraints on the diffusion and use of open source AI models, it will only slow our ability to win this technological race. To ensure continued American geopolitical leadership and safeguard national security, America must accelerate *both* the development of domestic AI open source models and their diffusion across the globe. At the same time, the United States needs to continue prioritizing the propagation of American-based industry standards and best practices in AI development and deployment. Longstanding policy deliverables, such as the NIST AI Risk Management Framework (AI RMF), have been effective mechanisms for promoting American AI policy abroad, influencing allies and partners to embrace more pro-innovation governance approaches to AI. It is imperative that any "AI Diplomacy" strategy continue to leverage the NIST AI RMF as a means of influencing AI policy alignment between the United States and global allies.

[4] Todd Young, "A Tech Power Playbook for Donald Trump 2.0", *National Interest*, 10 Feb. 2025, available at https://nationalinterest.org/blog/techland/a-tech-power-playbook-for-donald-trump-2-0.
[5] *Ibid*.

Recommendations

- Deputize the Office of the Special Envoy for Critical and Emerging Technologies as the lead agency tasked with developing a comprehensive global engagement strategy for enacting an "AI Diplomacy" agenda to support American AI interests and, importantly, to protect open source AI development.[6] Supporting open source AI diffusion should be a key cornerstone of such an agenda. However, additional features of this strategy should also include:
    - Protecting US market access and promoting an open-innovation-oriented approach globally, including by advocating for adherence to technical standards.
    - Pushing back against efforts that would undermine the ability of American companies to rely on flexible and reasonable IP frameworks, such as the fair use doctrine, for AI training and development.
- Instead of tamping down on the proliferation of open source AI, policymakers should instead highlight American open source AI leadership and make it a cornerstone of this new "AI diplomacy" priority in international engagements. The Trump Administration should leverage America's competitive advantage in this space and ensure the proliferation of policies and technologies that promote an open AI future. To balance the need for American open source AI diffusion while protecting vital technology, future export control restrictions and requirements should focus on hardware, not AI model weights or software.
    - Such controls could look to restrict the volume of GPUs target countries would be able to deploy running in parallel within a given data center and/or geographic region.
- The administration should additionally ensure that NIST continues to promote US industry-driven, consensus-based standards and frameworks (such as the NIST AI RMF) for AI to advance security and trust and enable global adoption.

Prioritize Risk-Based Approaches to Regulation

AI can be used in many ways across industries and the entire economy, which is why it is so important to focus any regulations on *use* and not the technology itself. AI is a tool, and like any tool it can be misused. But we do not regulate the underlying tools; we regulate how they are used in different contexts. AI should be no different.

[6] Legislation like the "Democracy Technology Partnership Act" could potentially serve as a template for a structure that would establish "a US interagency office at the State Department, tasked with creating a partnership among democratic countries to help set international standards and norms, conduct joint research, and coordinate export controls and investment screening on emerging and critical technologies." See https://www.warner.senate.gov/public/index.cfm/pressreleases?ID=3C5C0D90-479F-4657-A66D-93A1120B938E.

To encourage AI innovation and protect consumers, the greatest regulatory attention, resources and controls should be placed on the specific uses of AI that pose the greatest risk to people and their wellbeing. Such controls could include:

- Voluntary impact assessments (not mandatory third-party audits) and bias testing for high-risk AI and transparency about such self-testing.

- Definitions of high-risk and prohibited use cases.

- Clear differentiations between obligations for developers and deployers.

- Disclosure and transparency requirements to make known when people are interacting with AI and what datasets were used to train the AI model.

When determining risk, policymakers should look at those situations in which AI is used to make a "consequential decision" – which should be defined as "a decision that has a material legal or similarly significant effect on the provision or denial to any consumer of specifically enumerated appropriate high-risk use cases to promote harmonization and reduce burden on companies" – that could negatively impact an individual's fundamental rights. Importantly, many of the situations in which decisions may be rendered that impact such rights are already well-established tenets of existing law and regulation (e.g., protected class characteristics). The administration should focus regulatory efforts *only* on those areas that are currently unaddressed by existing statutory authorities. Where gaps exist, such as in requirements for transparency and documentation governance, the administration should consider *narrow and targeted action*.

Perhaps most importantly, budding interest in AI policy has resulted in a deluge of state legislative activity in this space. Many of these bills, while well-intentioned, are poised to wreak havoc on America's domestic AI industry. A lack of harmonized definitions, clarity in proposed provisions, and conflicting obligations are creating a budding patchwork of AI laws that could threaten America's leadership in this technology. As some of these bills could come into effect as soon as 2026, the Trump Administration should make it a near-term priority to work with Congress on a federal AI bill that preempts state AI legislation.

Recommendations

- The Trump Administration should work with Congress to prioritize legislation that preempts a patchwork of state AI legislation. Such legislation could also include reasonable guardrails in areas not addressed by existing law and regulations, such as:

    - Incentivizing "transparency governance" in the development and deployment of AI systems – in particular, transparency governance for training data sets and best practices for data provenance and data lineage – while recognizing the distinct roles of different actors throughout the AI supply chain. Understanding the source of data and the way it makes its way through the data pipeline to be used for AI model training

is a superior mechanism for enhancing trust and confidence in this technology, rather than policies that would seek restrictions to accessing model weights.

- More generally, we recommend the Trump Administration's AI Action Plan focus thematically on regulating the use of AI, based on risk, regardless of whether the underlying AI model is open or closed. Novel regulatory approaches may be warranted at some point to address potential risks of foundation models, but the consequences of an imprecise approach to regulation are too great to risk untested solutions.

- Direct the Bureau of Industry and Security (BIS) at the Department of Commerce to amend and clarify the Biden Administration's AI Diffusion Rule on the use of License Exception Low Processing Performance (LPP) to make clear that the use of distributors and other intermediaries, as well as in-country transfers, *is permitted* when the exporter or re-exporter has specific knowledge of the ultimate consignee of the controlled items being exported.

Rapidly Scale AI Usage to Streamline Government Operations

During the first Trump Administration, the Office of Management and Budget (OMB) recognized the urgent need for cross government solutions to "improve the ability of agencies to deliver mission outcomes, provide improved services, and effectively steward taxpayer dollars." To operationalize the vision, then Acting OMB Director Russel Vought issued OMB Memorandum M-19-16, on Centralized Mission Support Capabilities for the Federal Government. In short, this forward-looking guidance refreshed the entire government's approach to shared services. The vision of M-19-16 was to sidestep modernization of dozens of outdated, bespoke legacy IT systems, and simply replace them with high performing cross government platforms.

Recognizing its value, the subsequent Administration doubled down on this approach and embraced commercially hosted shared services platforms that deliver enormous capability in areas like human resources, payment processing, and claims processing. Most of these managed services are hosted in commercial clouds that are fully compliant with government information security standards and operated by trusted private sector partners. They are highly cost-effective because the entire burden for updates and infrastructure is born by the commercial provider. In many cases, machine learning, process automation, and generative AI have already been integrated into these systems. Government agencies have access to a strong ecosystem of approved providers via a cost efficient as-a-service model.

Because the government's data, business processes, and agency functions are already embedded in these platforms, and because they are built on best-in-class commercial software, they are primed for further scaling and integrating AI into nearly every facet of government operations. As new AI models emerge and the technology evolves, these advances can easily be integrated into government operations via these platforms.

Recommendations

- The Trump Administration should build on the success of the 2019 OMB Memo M-19-16 by designating proven, low-cost, high-efficiency shared service platforms as the preferred avenue to inject AI into existing government business functions. This closely aligns with the vision that OMB Director Vought outlined in 2019, when business process automation and machine learning were just beginning to be integrated into federal agency IT systems.

IBM recommends the new Administration focus on five areas with immediate streamlining benefits:

- Use AI to target fraud, waste, and abuse by implementing AI-driven analytics across key agencies to identify irregularities, predict potential risks, and enhance fraud detection.

- Use AI to transform how IT systems are modernized by using commercially hosted shared services platforms to automate, streamline, and accelerate the IT and application modernization process. AI can help accelerate the process of shedding costly to maintain legacy systems. This will reduce operational costs.

- Prioritize the development and deployment of AI-powered cybersecurity solutions to build adaptive systems and services that identify and mitigate cyber threats.

- Unlock human resources productivity by implementing a cross-agency, industry-hosted shared service platform that centralizes tools such as AI-enabled assistants for employee self-service and automates human resources tasks such as payroll, benefits administration, time entry, and performance management.

- Ruthlessly eliminate manual processes by adopting AI-powered automation tools. Automating manual tasks like data entry, form processing, and streamlining workflows allows for real-time insights, and enables faster, more informed decision-making.

Conclusion

We appreciate the opportunity to comment on this RFI and are likewise appreciative of the administration's strong commitment to ensuring continued American AI leadership. We look forward to continuing to work with NSF, OSTP, and other agencies to advancing a shared vision of a Golden Age of American technological innovation.