March 15, 2025

Faisal D'Souza, NCO
AI Action Plan
2415 Eisenhower Avenue
Alexandria, VA 22314

Re: AI Action Plan

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.

VIA EMAIL: [redacted]

Dear Mr. D'Souza,

Cloudflare welcomes the opportunity to submit input for the Development of an Artificial Intelligence (AI) Action Plan.

The AI Action Plan has the potential to shape the future of AI artificial intelligence in the United States and ensure that we maintain our competitive advantage in the field. We write to encourage you to address five issues in the AI Action Plan: (1) fostering conditions for US leadership in AI, (2) open-source AI; (3) distinguishing between AI model training and AI inference; (4) export controls; and (5) revision to the FedRAMP program.

I. Background on Cloudflare

Cloudflare is a leading connectivity cloud company that runs one of the world's largest networks, providing security, performance, and reliability services to millions of Internet properties. AI helps to make our performance and security

products smarter. Cloudflare powers Internet requests for more than 35% of the Fortune 500, and serves over 71 million HTTP requests per second on average. Cloudflare interconnects with approximately 13,000 networks globally, including major Internet Service Providers (ISPs), cloud services, and enterprises.

Cloudflare offers a wide variety of services to improve the security, reliability, and performance of websites, networks, and other applications online. Cloudflare's security services work by directing traffic to Cloudflare's network rather than directly to a hosting provider or internal network. Cloudflare then uses its "points of presence" in more than 335 cities in over 125 countries to screen traffic for cybersecurity risks and to cache content at the network edge in order to improve reliability and performance.

Cloudflare stands at the forefront of AI innovation, leveraging its global network and serverless architecture to drive efficiency, enhance security, and redefine the economics of AI computing. AI has been a foundational element of Cloudflare's security and performance offerings. Since its inception, Cloudflare has functioned as an AI-powered security platform, analyzing vast amounts of Internet traffic to detect and mitigate cyber threats in real-time. Our machine learning-based security systems continuously identify and neutralize emerging threats before they are recognized by traditional security measures. This proactive approach ensures that our customers benefit from smarter, more adaptive security solutions.

One of the most transformative applications of Cloudflare's AI capabilities is within our Workers AI platform, which allows developers to run open-source AI models in a serverless way and has emerged as a game-changer for AI inference. Developers are increasingly leveraging Workers AI to build AI-driven applications, particularly AI agents, due to its cost-efficient, pay-per-use model. Cloudflare's serverless inference model dynamically allocates processing power based on real-time demand, resulting in significant cost savings for customers. Cloudflare's AI Gateway further enhances efficiency by allowing developers to cache AI-generated responses, significantly reducing latency and improving performance. AI Gateway is designed to help developers and organizations building AI applications better monitor, control, and optimize their AI usage. Developers can serve AI requests

directly from Cloudflare's cache rather than relying solely on model providers. This innovation is driving the next wave of AI-powered applications.

Cloudflare sits in a unique position in the AI ecosystem, with customers that include both a significant portion of the Internet's content creators and many of the world's largest AI companies. Sitting between those two different types of entities puts Cloudflare in an important position to help determine the technical mechanisms of how AI crawlers and agents are allowed, and on what terms, and how the AI-driven web of the future will fit together. Cloudflare has already rolled out tools for content creators that provide visibility into how AI crawlers crawl their sites, as well as mechanisms to block AI crawlers.

II. Fostering the conditions for U.S. leadership in AI

For AI to reach its full potential in the United States, the government must establish the right preconditions to foster innovation, collaboration, and responsible development.

AI systems thrive on access to diverse, high-quality datasets. At the same time, AI systems often destroy the value proposition for building those data sets. For example, AI crawlers that scan the web looking for content to train LLMs often fail to drive traffic to the sites they scan. Without the ability to control scanning and realize value, site owners will be discouraged to launch or maintain Internet websites. Creators are more likely to place their content behind paywalls, with the largest publishers trying to negotiate direct deals. AI model providers will in turn struggle to find and access the long tail of high-quality content on smaller sites. To avoid this outcome, we will need an environment that enables a transparent exchange of permissions and value. Industry has begun to work on new standards and tools to address this concern and maintain a healthy market for content. We believe the United States can serve as a leader in encouraging this industry-led approach, preventing overly restrictive policies on data movement or access that can hinder technological advancements, limit competition, and put U.S. companies at a disadvantage globally.

Interoperability is another critical pillar for AI growth. AI systems and services must be designed to work well together, allowing businesses and developers to leverage multiple models, platforms, and cloud providers. Promoting open standards is key to preventing AI silos, reducing inefficiencies, and fostering a competitive, dynamic ecosystem. For example, Cloudflare's Workers AI platform is model-agnostic, supporting a multi-model approach that allows companies to build AI pipelines across different models seamlessly. Just as multi-cloud strategies help businesses avoid vendor lock-in and distribute workloads efficiently, an interoperable AI ecosystem will enable greater flexibility and innovation across industries. Without these principles, restrictive policies could push businesses toward more innovation-friendly markets abroad, weakening U.S. leadership in AI.

Interconnection and the availability of low latency edge networks also play a fundamental role in AI's infrastructure. The more networks are interconnected on the Internet, the more resilient AI-powered services become. Fast, low-latency AI applications depend on efficient routing and direct network connections, which require robust peering relationships. Currently, more than 99% of global peering agreements are settlement-free, allowing networks to interconnect without additional costs, which benefits end-users and businesses alike. Cloudflare, for example, peers with over 12,500 other networks, ensuring that users around the world can seamlessly and efficiently access AI applications close to where they are located. Encouraging policies that promote network interconnection and neutral peering arrangements will strengthen AI infrastructure, ensuring that AI services can scale rapidly and perform efficiently across different environments.

AI is rapidly evolving into a foundational technology, much like electricity in the early 20th century—transforming industries, boosting productivity, and enabling new forms of creativity. AI-powered tools already streamline software development, enhance content creation, and support medical advancements. However, its continued success depends on a strong, interconnected digital infrastructure that supports real-time data processing and secure cross-platform collaboration. The U.S. government must recognize that AI is not developing in isolation—its growth is deeply tied to how networks, data flows, and

interoperability standards are managed. By fostering an open, interconnected, and data-driven AI ecosystem, the U.S. can secure its leadership in AI, drive economic growth, and ensure that AI technologies align with democratic values and global competitiveness.

III. The AI Action Plan should endorse open-source AI

Open-source AI is pivotal for the United States to maintain its leadership in technological innovation, enhance national security, and foster economic growth. Cloudflare strongly encourages the United States to endorse the availability of open-source AI in the AI Action Plan. Open-source AI promotes transparency, collaboration, and rapid advancement by allowing researchers, developers, and organizations to access, modify, and improve AI models collectively. Access to open-source AI enables developers to build new applications and create new businesses, without having to rely on a small number of gatekeepers.

To maintain U.S. leadership in AI, the government should permit the free global transfer of open-source models. Open ecosystems drive innovation—broader adoption leads to better tools, optimizations, and integrations. If U.S. models are restricted, foreign alternatives will dominate, shaping AI technology around adversarial interests.

History shows that closed systems stagnate while open ones thrive—Linux succeeded where proprietary Unix failed. The same applies to AI. Limiting U.S. models risks ceding influence and weakening America's role in AI's future. To stay ahead, the United States must ensure its open-source AI models become the global standard.

Open-source AI models have been adopted by a diverse community, including U.S. government agencies working on defense and national security applications, as well as private sector partners. This collaborative approach accelerates AI development and ensures that advancements are aligned with national interests. The Cybersecurity and Infrastructure Security Agency (CISA), for example, emphasizes

that developers of AI models can learn from the open-source software community to enhance security and resilience. Similarly, the Commerce Department's Bureau of Industry Security (BIS) has also traditionally recognized the value of open-source technologies by creating special decontrols which reduce the export control barriers on companies seeking to use open-source code. By adopting best practices from open-source development, AI systems can become more robust and secure, mitigating potential risks associated with AI deployment.

By endorsing open-source AI, the U.S. government can harness collective expertise, drive innovation, and ensure that AI technologies are developed in a manner that upholds democratic values and national security interests. This commitment would position the United States at the forefront of the AI revolution, fostering an ecosystem where transparency, collaboration, and security are paramount.

IV. The AI Action Plan should distinguish between the infrastructure required for training AI models and AI inference

It is critical that the AI Action Plan clearly distinguish between the infrastructure required for the training of AI models and the infrastructure required for AI inference tasks. While the training of AI models and AI inference are essential, and interdependent, components of the AI ecosystem, they serve distinct functions and involve different levels of risk. They also serve distinct roles in the lifecycle of AI systems and require tailored policy approaches to address their unique challenges and opportunities.

AI models are developed through a computationally intensive process called training. During this phase, vast datasets are used to teach the model to recognize patterns and relationships in data. For example, a facial recognition model might be trained on millions of images to identify key features like eye shape or hair color. This training process requires significant computational resources, specialized hardware (e.g., GPUs or TPUs), and large amounts of energy, making it both costly and resource-intensive. The same high-performance GPUs used for training large language models can also be leveraged for building military AI applications (e.g.,

autonomous drone navigation, target recognition, and homing capabilities), raising the risks of access to large-scale GPU clusters capable of supporting such uses.

In contrast, AI inference is the operational phase where a trained model applies its learned knowledge to new, unseen data to generate predictions or outputs. For instance, a self-driving car recognizing a stop sign it has never encountered before exemplifies AI inference in action. Unlike training, inference often occurs in real-time and is deployed across diverse environments, from edge devices like smartphones to cloud-based systems. This phase requires policies that focus on ensuring AI application developers have access to AI inference infrastructure that is optimized with low-latency performance, cybersecurity safeguards against adversarial attacks, and privacy protections for end users whose data is processed by these models. U.S. companies providing AI inference infrastructure services should be allowed to operate globally with minimal restrictions, as limited GPU deployments for inference alone cannot feasibly train models posing national security risks. Enabling these companies to become the global providers of choice directly supports U,S. strategic interests. Conversely, restricting their ability to do business will inevitably cede market dominance to foreign competitors, including adversaries.

The U.S. government should, therefore, avoid imposing unnecessary restrictions on U.S. companies offering AI inference infrastructure globally. Encouraging openness will drive growth and solidify U.S. technological dominance by making its infrastructure the preferred choice worldwide. On the other hand, limiting U.S. providers will create a gap that foreign competitors—including adversaries—will inevitably fill.

Policies that conflate the training of AI models with AI inference tasks risk imposing undue regulatory burdens on inference infrastructure providers. Ensuring that regulations do not inadvertently stifle AI adoption and deployment at the inference level is crucial for fostering economic growth and maintaining U.S. leadership in AI development. By clearly differentiating between the infrastructure required to train AI models and the infrastructure required for AI inference tasks, the AI Action Plan

can establish more effective and proportionate policies. These respective targeted policies will help to foster innovation while ensuring U.S. leadership in AI. A well-balanced policy framework that acknowledges these differences will allow the U.S. to promote AI innovation while addressing legitimate concerns around AI risks and misuse.

V. The AI Action Plan should support export control exemptions for equipment shipped abroad for U.S. company use, ensuring continued U.S. leadership in AI.

In addition to endorsing open-source AI models, the U.S. government should also continue to ensure that U.S. companies have access to the hardware and other equipment needed to support AI research and development. In order to maintain U.S. leadership in AI, U.S. companies need access to advanced computing / supercomputing systems, such as advanced computing integrated circuits (ICs), to train AI models and use those AI models to perform inference tasks. While we recognize there are legitimate national security reasons for why the U.S. government will have an interest in regulating the exports of large quantities of advanced computing ICs, these controls should be narrowly tailored to the risks they are seeking to mitigate. In the case of AI innovation, the training of new AI models by U.S. adversaries, which requires intensive computation in a concentrated area, presents significantly more national security risks than the use of much smaller numbers of advanced computing ICs to deploy AI models.

Many U.S. companies are becoming leading innovators in the development of Software-as-a-Service products that incorporate AI, but in order to maintain their competitiveness they need to be able to deploy their products around the world, close to their users. The U.S. government should continue to look for ways to implement controls that are flexible and account for the needs of U.S. companies to export advanced computing ICs and associated hardware outside the U.S. For example, license exemptions for U.S. companies to export hardware up to a certain threshold of total processing power (TPP) for their own use in research and

development should always be considered when implementing any new controls on exports of advanced computing ICs.

VI. The AI Action Plan should optimize FedRAMP to streamline AI procurement within the federal government and strengthen U.S. competitiveness

The Federal Risk and Authorization Management Program (FedRAMP) is a crucial framework for securing cloud services used by U.S. government agencies. However, its current structure is often slow, costly, and burdensome, limiting the government's ability to rapidly procure and deploy cutting-edge AI solutions. As AI innovation accelerates globally, optimizing and modernizing FedRAMP is essential to ensure that U.S. agencies can leverage AI-driven tools efficiently while strengthening the domestic AI ecosystem.

The AI Action Plan should prioritize reforms that enhance the efficiency, scalability, and responsiveness of FedRAMP, enabling the inclusion of more AI tools in the FedRAMP marketplace. Modernizing FedRAMP would not only improve the U.S. government's ability to adopt AI but would also strengthen U.S. technology leadership. A more agile and transparent authorization process would encourage more AI and cloud companies—particularly startups and small businesses—to pursue government contracts, fostering a more competitive marketplace. By making FedRAMP more efficient, the government can accelerate its adoption of AI while ensuring that U.S. companies remain at the forefront of global AI innovation.

VII. CONCLUSION

Cloudflare appreciates the opportunity to contribute to the development of the AI Action Plan and urges the U.S. government to adopt policies that foster AI innovation, strengthen national security, and maintain global leadership. By supporting open-source AI, ensuring clear distinctions between AI training and inference, refining export controls, and modernizing FedRAMP, the United States can create a policy framework that promotes competition, protects critical

infrastructure, and enables responsible AI development. Cloudflare stands ready to collaborate with policymakers to shape an AI ecosystem that is open, secure, and globally competitive.

Sincerely,

/s/ Zaid A. Zaid

Zaid A. Zaid
Director, Head of U.S. Public Policy