moz://a

Mozilla's Comments on The Establishment of Reporting Requirements for the Development of Advanced Artificial Intelligence Models and Computing Clusters

About moz://a

Mozilla's mission is to ensure the internet is a global public resource, open and accessible to all. An internet that truly puts people first, where individuals can shape their own experience and are empowered, safe, and independent.

Founded as a community open source project in 1998, Mozilla consists of several organizations, most notably the non-profit Mozilla Foundation, which leads our movement-building work, and its wholly owned subsidiary, the Mozilla Corporation, which leads our market-based work, including the development of the Firefox web browser. They work in close concert with each other and a global community of tens of thousands of volunteers under the single banner: Mozilla.

For the past five years, Mozilla has been committed to advancing trustworthy AI. Mozilla published a paper in early 2024, Accelerating Progress Toward Trustworthy AI, that outlines how Mozilla and its allies are advancing openness, competition, and accountability in AI. Mozilla is putting its resources behind these priorities as well: The Mozilla Foundation has been dedicating 100% of its $30M a year budget to philanthropic activities, advocacy, and programmatic work on this topic. Mozilla is also investing another $30M in research and development on trustworthy AI via Mozilla.ai, as well as $35M in responsible tech startups — including startups with a focus on trustworthy AI — through Mozilla Ventures and the Mozilla Builders accelerator program. On the frontlines of modern AI practices, Mozilla freely provides an open-source, large-language model (LLM) AI model deployment system for local use and empowers more people to enhance the safety of models online through the 0Din bug bounty program. In addition, Mozilla supports the work of academic and civil society organizations engaged in artificial intelligence research, including the Open Source Initiative's work on defining Open Source AI and through the work of the Columbia Convening which brought together a diverse set of stakeholders to talk about Openness and AI.

As an independent and mission-driven organization, Mozilla is committed to working with regulators to develop effective policies that ensure that innovation and growth in AI serve the public interest. In the past, Mozilla has provided comments on critical topics like NTIA's consultation on openness in AI models, NIST's AI Risk Management Framework, and NIST's Request for Comments related to Managing Misuse Risk for Dual-Use Foundation Models.

Executive Summary

Mozilla has been on the frontline of defending the open internet for 25 years. Our history as an organization is deeply intertwined with that of the open source movement, and we believe that this same openness is critical for AI as its adoption and development increase exponentially. Not only does openness lead to safer software that is less vulnerable, it helps with mitigating bias and harnesses the "wisdom of the crowd" to move technological frontiers forward while providing significant economic benefits to society.

We appreciate the effort that BIS has invested into developing updated reporting requirements for the development of advanced AI models and computing clusters and its willingness to consult with external stakeholders on how to best implement and update the requirements. Mozilla is pleased to offer its feedback and ideas on ensuring that any updates are both successful and broadly applicable for the open source community rather than the largest big tech research and development entities. We hope that BIS will further consider the needs and value of the open source community as it drafts requirements, especially given the important role which open source is playing in the AI ecosystem and U.S. economy.

The following list highlights key feedback from this document:

1. When updating the proposed rule, BIS should take into account and provide details as to how such rules would apply to the open source community given the unique structure of many such open source projects, which may lack a specific owner or controlling entity and be geographically distributed and freely available around the world. Ideally, BIS would work in consultation with civil society organizations such as the Open Source Initiative (OSI) on developing a framework which fits the characteristics of the open source AI community.
2. AS BIS continues to update relevant technical conditions related to collection thresholds detailed in the proposed rule based on technological advances, we ask that BIS consider establishing a minimum update cycle of six months given the rapid pace of change in the AI landscape. This is necessary to maintain BIS' core focus on the regulation of frontier models and to not unnecessarily stymie innovation across the broader AI ecosystem.
3. BIS should provide additional definitional clarity as to what 'planned applicable activities' refers to and when a project or idea meets the definition of having been 'planned.'

Our Feedback

Taking Open Source Into Account

We recommend that as BIS further develops the proposed rule, it should take into account the unique attributes of the open source AI ecosystem in order to more effectively provide guidance to a large swath of the broader AI ecosystem. Open source AI is a rapidly developing field, and many AI models are launched at varying degrees of openness, which may include those who comply with

the full definition of open source as developed by OSI or whose AI models are only partially open, making it critical for BIS to provide clarity focused on the unique attributes of the open source ecosystem. For example, what are the obligations of an American software developer who contributes to a large open source AI project in their free time that eventually falls under BIS' proposed collection thresholds? In this scenario, the developer may have little to no insight into who began the project or their physical location, plans for the project in the future, and likely has no explicit control over the project. Providing guidance for such common scenarios in the open source AI community would be helpful in enhancing adherence to the rule and creating clarity for many individuals, nonprofits, and research institutions which may be relatively unfamiliar with BIS and not have the resources to analyze proposed regulations that may affect them.

It is critical that the Open Source AI community is proactively considered during the drafting and promulgation of important regulations given the positive economic and national security impact of open source on America's global competitiveness. As highlighted in Mozilla's response to NTIA, from Google's decision to openly document and share its newly developed transformer architecture which underpins all state-of-the-art large language models to the open-sourcing of libraries for techniques like "Low-Rank Adaptation" (LoRA), open-source has helped to propel the AI frontier forward. Initiatives like Mozilla's Common Voice project, which is working to compile the world's largest crowdsourced multilingual dataset, has helped to lower barriers of entry in fields like voice technology while also expanding access for communities frequently underrepresented in machine learning training data. According to a 2024 Harvard Business School Working Paper, open source software provides economic value estimated above $8 trillion dollars thanks to its near ubiquitous use across online applications that, if lost, would necessitate huge investment and costs by firms, government, or others to replace. The economic case for open source software is clear, and as the use of AI and open source AI models in particular continues to rise, it is important that BIS and other agencies appropriately weigh any regulatory actions which may cause harm to the open source community with the commensurate cost to America's economic security and proactively consider the open source AI community when drafting and promulgating rules.

Knowing exactly for whom proposed rules are applicable is critical in not inadvertently creating a chilling effect on the broader open source AI ecosystem, especially given the spectrum of openness which many AI models operate under today. Mozilla recommends that BIS consult with non-governmental organizations such as the Open Source Initiative which has produced the most widely accepted definition of Open Source AI in consultation with a broad spectrum of stakeholders. While not directly discussed in the proposed rule, Mozilla would like to surface previous comments provided to NTIA and those highlighted in a joint letter from Mozilla and CDT to Secretary Raimondo related to the risk of leveraging export controls on open-source AI. Using such a blunt instrument may lead to unintended economic security consequences by obstructing innovation, research, competition, and transparency, while simultaneously failing to meet national security objectives due to significant practical obstacles to effective implementation and enforcement of such actions.

Ongoing Collection Threshold Updates & Survey Questions

BIS' work updating technical conditions to more appropriately address the frontier of AI models is commendable and will be beneficial in minimizing the regulatory burden from the rule for smaller entities. However, as the field of Artificial Intelligence is moving extraordinarily rapidly, with advances happening nearly daily, it is critical that BIS engages in an ongoing process of updating collection thresholds. As directed by section 4.2(b) of E.O 14110, "BIS will update these technical conditions as appropriate." BIS can improve the update process by further defining what "appropriate" means in such a rapidly changing environment. One such way to do so would be to create a minimum time period after which BIS will engage in a review and update of the relevant technical conditions, such as doing so every six months. Doing so will help to ensure that BIS' focus remains on frontier models rather than inadvertently capturing a larger swath of AI model providers unnecessarily due to a lack of definitional updates completed in a timely manner.

As BIS considers future questions for relevant firms as part of its ongoing monitoring work, we ask that BIS takes into account broader goals like transparency and accountability related to frontier AI products and solutions when deciding upon questions to include in the survey and requests for data. This can include specific questions on topics such as the data used to train the AI model and the results of red-team testing related to bias, security vulnerabilities, and other concrete harms. If such data can be made available through this process, it will help to inform future government efforts when tackling challenges related to AI's impact on the economy and society.

Definitional Clarity

The proposed rule would require that "Covered U.S. persons subject to the reporting requirements in paragraph (a)(1) of this section must notify BIS of 'applicable activities' via email each quarter, identifying any 'applicable activities' planned in the six months following notification." We ask that BIS provide additional definitional clarity related to "planned applicable activities," specifically, when a project or idea would meet the definition of having been "planned." One could see a scenario where an entity may not think to inform BIS of such a "planned applicable activity" when it is in the early stages of internal development within a startup, is a plan without the requisite resources to actually implement, etc. Providing additional clarity is likely to increase adherence to the rule and minimize unnecessary notifications to BIS which would create additional and unnecessary burdens for both BIS and complying entities.

Increasing Cyber Resilience

According to the proposed rule, "the U.S. Government must minimize the vulnerability of dual-use foundation models to cyberattacks... Accordingly, the U.S. Government needs information about the cybersecurity measures that companies developing dual-use foundation models use to protect

those models, as well as information about those companies' cybersecurity resources and practices." This comment is clearly tailored to companies with private code bases and it would be helpful if BIS also provided clarification as to how their thinking pertains to open source projects where a spectrum of AI artifacts such as model weights and datasets may be publicly available. In many situations, increased openness can enable additional security and decrease vulnerabilities: the use of bug bounty programs and public red-teaming are a testament to the benefit of more public scrutiny on code bases and models, hence their widespread adoption and support from the institutions as diverse as the Defense Digital Service's Hack the Pentagon initiative to the White House's support for a "red-teaming" event focused on large language models to myriad private companies, including many leading AI labs who have voluntarily agreed to engage in external security testing of AI systems before their release. As highlighted in Mozilla's March 2024 comments to NTIA, research drawing on open source AI has helped to advance red-teaming and safety alignment work, removing protections from (or 'jailbreaking') aligned models (including state-of the art proprietary models). This has not only helped to advance AI safety and security research and make open source AI safer, but enabled providers of proprietary AI models to address critical product vulnerabilities.

In order to effectively mitigate cyber threats which may have an impact on relevant AI models, or on any AI models which may have an impact on national security, BIS should consider how to effectively share information gathered about cyber threats and vulnerabilities with both covered entities under the proposed rule as well as other AI companies and open source AI communities. Given that more than 90% of commercial programs use open source software, it seems likely that many entities, including government agencies and leading AI labs producing frontier models, will depend on open source for their work. By proactively creating a strategy for sharing threat intelligence and known vulnerabilities with the broader AI community including open source providers, BIS could further its national security objectives. Mozilla has a long history of developing bug bounty programs and proliferating security standards across the web, including with the recently launched 0din – a bug bounty program focused on large language models – and stands ready to support the government's work to share threat intelligence information across the open source AI ecosystem.

Open Source AI, National Security, & Economic Security

Open source can have significant benefits for national security by increasing successful adoption of AI in critical applications and sectors, bolstering resilience by reducing dependencies on a small number of companies, helping expose vulnerabilities early, and gathering feedback from much broader audiences. According to a 2024 CSIS Preliminary Assessment on Defense Priorities in the Open-Source AI Debate, "Open-source software and standards are already widespread in U.S. national security applications… Open-source software is ubiquitous, permeating over 96 percent of civil and military codebases, and will remain a core piece of defense infrastructure for years to come." The paper further discusses national security benefits from open-source AI models, saying

"More critically, the existence of open foundation models mitigate dependence on single vendors when sustaining AI-powered defense systems," and that "Preliminary evidence suggests that open foundation models might benefit the defense industrial base."

The emerging discussion around Public AI has also highlighted the national security benefits of open and public AI, with Vanderbilt Policy Accelerator's 2024 report titled The National Security Case for Public AI also highlighting the dangers of "Reliance solely on unregulated national champions makes the U.S. government dependent on a small number of firms-and even on particular individuals. This is a tactical and strategic national security risk because one person or firm holds considerable power over the government." Both Public AI and open source AI (which are deeply intertwined) help to create more optionality for the government, preventing vendor lock-in and providing national and economic security benefits. It may behoove BIS to conduct a cost-benefit analysis in consultation with the open source AI community, government software providers, and government procurement experts related to open source AI's potential benefits and risks related to national security, taking into account benefits such as mitigating vendor lock-in.

Final Thoughts

Mozilla would like to once again thank BIS for its work in developing new thresholds for data collection, working to minimize the potential negative impact of such rules on a large swath of the AI community, and in making an effort to effectively balance the benefits and potential risks of AI with regards to national and economic security. The proposed rule lays out thoughtful questions on how BIS can optimize its work related to frontier model data collection further.

Mozilla believes that the proposed rule and BIS' larger efforts would be helped with an increased focus on the needs of the open source community, which plays a large role in pushing the AI frontier forward and which has demonstrated clear benefits for the United States through increased transparency, accountability, and optionality for end users. Due to the pace of change in the AI industry, BIS should also consider creating minimum update cycles for its collection thresholds in order to not inadvertently collect more information than necessary and create undue regulatory burdens. In addition, BIS can provide definitional clarity on certain terms such as 'planned applicable activities' to further enhance adherence to the proposed rule and to minimize additional paperwork burdens for both BIS and companies who may incorrectly believe they have to comply with such rules. Finally, BIS' proposed rule discusses in some depth the potential national security risks emanating from advanced AI models and related cyber risks. In order to mitigate such risks, BIS should further examine how to effectively share information on vulnerabilities and cyber threats with both frontier AI companies and the broader AI ecosystem, especially open source entities.

We hope that BIS further considers the potential impact of the proposed rule and future regulatory actions on the open source community and appropriately weighs the myriad benefits which open

source AI and open source software more broadly produce for America's national and economic security. Mozilla thanks BIS for the opportunity to provide our feedback and we look forward to working with BIS further in the future.