

March 13, 2025

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

From: Christopher Lehane
OpenAI
1455 3rd Street
San Francisco, CA 94158

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.

“It is the policy of the United States to sustain and enhance America’s global dominance in order to promote human flourishing, economic competitiveness, and national security”
– President Donald J. Trump, Executive Order 14179, January 23, 2025

OpenAI respectfully submits the enclosed proposals to the Office of Science and Technology Policy as it weighs a new AI Action Plan that will, as Vice President Vance stated recently at the Paris AI Action Summit, maintain American leadership in AI and “make people more productive, more prosperous, and more free.” As America’s world-leading AI sector approaches artificial general intelligence (AGI), with a Chinese Communist Party (CCP) determined to overtake us by 2030, the Trump Administration’s new AI Action Plan can ensure that American-led AI built on democratic principles continues to prevail over CCP-built autocratic, authoritarian AI.

OpenAI agrees with the Trump Administration that AI creates prosperity and freedom worth fighting for—especially for younger generations whose future will be shaped by how this Administration approaches AI. Globally, most ChatGPT users are under age 35; in the US, about one third are ages 18 to 24.¹ Both young people and their parents recognize the economic opportunities AI presents:

- More than seven in 10 parents in the US believe children today will be worse off financially than they are.²
- Nine in 10 US parents think it's important that their kids learn how to use artificial intelligence for their future jobs—and eight in 10 say either that isn’t happening today, or they don’t know if it is.³
- Three in four college-age AI users want to use AI in their education and careers. Many are teaching themselves and their friends about AI without waiting for their schools to provide formal AI education.⁴

¹ Self-reported among logged-in users

² Pew Research Center: [Views of children’s financial future](#), Jan. 2025

³ Morning Consult survey commissioned by Samsung: [88% of US Parents of Gen Alpha & Gen Z Students Say AI Will Be Crucial to Their Child’s Future Success](#), Sept. 2024

⁴ OpenAI: [Building an AI-Ready Workforce: A Look at ChatGPT Adoption in the US](#), Feb. 2025

OpenAI

In particular, AI could drive significant increased productivity over the next decade. Here's how we can realize this heightened prosperity and greater freedom together.

Scaling human ingenuity

Innovation creates and scales our ability to push beyond our current limits. From foot travel to the domesticated horse, the wheel, steam power, the car, the plane—we scaled the freedom of mobility. From daylight to candle and lamplight, to electricity providing light and power at all hours—we scaled the freedom to produce, think and create. From word of mouth to the stylus and tablet, to the printing press, telegraph, phone, computer, smartphone—we scaled freedom of learning and knowledge. Now, as we approach AGI, innovation is poised to scale human ingenuity itself—the sum of our freedoms to learn and know, think, create and produce.

As our CEO Sam Altman has written, we are at the doorstep of the next leap in prosperity: the Intelligence Age. But we must ensure that people have **freedom of intelligence**, by which we mean the freedom to access and benefit from AGI, protected from both autocratic powers that would take people's freedoms away, and layers of laws and bureaucracy that would prevent our realizing them.

More than 400 million people around the world are using ChatGPT to ideate, discover, and break through beyond what we're currently capable of doing on our own. Just two weeks ago, we partnered with the Department of Energy's national labs to bring together 1,500 scientists to use our tools to take scientific discovery farther, faster.

Our work at OpenAI also suggests that as AI advances, progress accelerates and becomes increasingly affordable, as reflected in these three scaling principles:

1. *The intelligence of an AI model roughly equals the log of the resources used to train and run it. Until recently, scaling progress has primarily come from training compute and data, but we have shown how to make intelligence scale from inference compute, as well. The scaling laws that predict these gains are incredibly precise over many orders of magnitude, so investing more in AI will continue to make it better and more capable. We believe that the socioeconomic value of linearly increasing intelligence is super-exponential in nature.*

2. *The cost to use a given level of AI capability falls by about 10x every 12 months, and lower prices lead to much more use. We saw this in the change in token cost between GPT-4 in early 2023 and GPT-4o in mid-2024, where the price per token dropped about 150x in that time period. Moore's Law predicted that the number of transistors on a microchip would double roughly every two years; the decrease in the cost of using AI is even more dramatic.*

OpenAI

3. *The amount of calendar time it takes to improve an AI model keeps decreasing.* AI models are catching up with human intelligence at an increasing rate. The typical time it takes for a computer to beat humans at a given benchmark has fallen from 20 years after the benchmark was introduced, to five years, and now to one to two years—and we see no reason why those advancements will stop in the near future.

By scaling human ingenuity ever faster and more affordably, AGI will create a flywheel of more freedom leading to more productivity, prosperity, and yet more innovation—letting us once again focus on positive-sum growth.

Advancing democratic AI

OpenAI believes the best future is one in which we move forward with democratic AI—AI that is shaped by the democratic principles America has always stood for. As OpenAI recently laid out in our [Economic Blueprint](#), we believe these principles include:

- *A free market* promoting free and fair competition that drives innovation.
- *Freedom for developers and users* to work with and direct our tools as they see fit, in exchange for following clear, common-sense technical standards that help keep AI safe for everyone, and being held accountable when they don't.
- *Preventing government use of AI tools to amass power and control their citizens*, or to threaten or coerce other states.

In advancing democratic AI, America is competing with a CCP determined to become the global leader by 2030. That's why the recent release of DeepSeek's R1 model is so noteworthy—not because of its capabilities (R1's reasoning capabilities, albeit impressive, are at best on par with several US models), but as a gauge of the state of this competition.

As with Huawei, there is significant risk in building on top of DeepSeek models in critical infrastructure and other high-risk use cases given the potential that DeepSeek could be compelled by the CCP to manipulate its models to cause harm. And because DeepSeek is simultaneously state-subsidized, state-controlled, and freely available, the cost to its users is their privacy and security, as DeepSeek faces requirements under Chinese law to comply with demands for user data and uses it to train more capable systems for the CCP's use. Their models also more willingly generate how-to's for illicit and harmful activities such as identity fraud and intellectual property theft, a reflection of how the CCP views violations of American IP rights as a feature, not a flaw.

OpenAI

Today, CCP-controlled China has a number of strategic advantages, including:

- *As an authoritarian state, its ability to quickly marshal resources*—data, energy, technical talent, and the enormous sums needed to build out its own domestic chip development capacity.
- *Its preexisting Belt and Road initiative.* As with Huawei, the PRC will scale the adoption of PRC-based AI systems like DeepSeek’s by coercing countries needing AI tools and nation-building infrastructure funds.
- *Its ability to benefit from regulatory arbitrage being created by individual American states* seeking to pass their own industry-wide laws, some of which are modeled on the European Union’s regulation of AI. These laws are easier to enforce with domestic AI companies than PRC-based companies and could impose burdensome compliance requirements that may hinder our economic competitiveness and undermine our national security. They also may weaken the quality and level of training data available to American entrepreneurs and the usefulness for downstream consumers and businesses.
- *Its ability to benefit from copyright arbitrage being created by democratic nations* that do not clearly protect AI training by statute, like the US, or that reduce the amount of training data through an opt-out regime for copyright holders, like the EU. The PRC is unlikely to respect the IP regimes of any of such nations for the training of its AI systems, but already likely has access to all the same data, putting American AI labs at a comparative disadvantage while gaining little in the way of protections for the original IP creators.

While America maintains a lead on AI today, DeepSeek shows that our lead is not wide and is narrowing. The AI Action Plan should ensure that American-led AI prevails over CCP-led AI, securing both American leadership on AI and a brighter future for all Americans.

What we propose

OpenAI’s freedom-focused policy proposals, taken together, can strengthen America’s lead on AI and in so doing, unlock economic growth, lock in American competitiveness, and protect our national security. Specifically, we detail:

A regulatory strategy that ensures the freedom to innovate: For innovation to truly create new freedoms, America’s builders, developers, and entrepreneurs—our nation’s greatest competitive advantage—must first have the freedom to innovate in the national interest. We propose a holistic approach that enables voluntary partnership between the federal government and the private sector, and neutralizes potential PRC benefit from American AI companies having to comply with overly burdensome state laws.

OpenAI

An export control strategy that exports democratic AI: For countries seeking access to American AI, we propose a strategy that would apply a commercial growth lens—both Total and Serviceable Addressable Markets—to proactively promote the global adoption of American AI systems and with them, the freedoms they create. At the same time, the strategy would use export controls to protect America’s AI lead, including by making updates to the AI diffusion rule.

A copyright strategy that promotes the freedom to learn: America’s robust, balanced intellectual property system has long been key to our global leadership on innovation. We propose a copyright strategy that would extend the system’s role into the Intelligence Age by protecting the rights and interests of content creators while also protecting America’s AI leadership and national security. The federal government can both secure Americans’ freedom to learn from AI, and avoid forfeiting our AI lead to the PRC by preserving American AI models’ ability to learn from copyrighted material.

A strategy to seize the infrastructure opportunity to drive growth: Sustaining America’s lead on AI means building the necessary infrastructure to compete with the PRC and its commandeered resources. We propose policies to seize this unmissable opportunity to catalyze a reindustrialization across our country, creating and supporting hundreds of thousands of jobs, boosting local economies, modernizing our energy grid, and preparing an AI-ready workforce—the key pillar of any country’s AI infrastructure.

An ambitious government adoption strategy: Advancing democratic AI around the world starts with ensuring that the US government itself sets an example of governments using AI to keep their people safe, prosperous, and free. With the PRC progressing toward ambitious targets for AI adoption across its public administration, security, and military, the US government should modernize its processes to safely deploy frontier AI tools at the pace of the private sector and with the efficiency Americans deserve.

America always succeeds when it bets on American ingenuity. The enclosed policy proposals are either derived from, or in the case of copyright represent updates to [OpenAI's Economic Blueprint](#), and we look forward to discussing them with you.



Chris Lehane
Vice President, Global Affairs

OpenAI

1. Preemption: Ensuring the Freedom to Innovate

We propose creating a tightly-scoped framework for voluntary partnership between the federal government and the private sector to protect and strengthen American national security. This framework would extend the tradition of government receiving learnings and access, where appropriate, in exchange for providing the private sector relief from the 781 and counting proposed AI-related bills already introduced this year in US states. This patchwork of regulations risks bogging down innovation and, in the case of AI, undermining America's leadership position.

Overseen by the US Department of Commerce and in coordination with the AI Czar, perhaps by reimagining the US AI Safety Institute, this effort would provide domestic AI companies with a single, efficient "front door" to the federal government that would coordinate expertise across the entire national security and economic competitiveness communities.⁵

This targeted framework would empower the federal government to:

- Work with both large AI companies and start-ups on a purely voluntary and optional basis to stay informed about AI risks as well as cutting-edge capabilities that support US national interests, including by establishing sandbox and testing capabilities on the secure premises of federal agencies.
- Evaluate the state of American AI technology against the technology of competitors and adversaries, including evaluating foreign models for the potential for back doors or malign influence.
- Coordinate the development of technical standards for evaluating and safeguarding frontier models from national security risks.
- Provide American AI companies with the tools and classified threat intelligence to mitigate national security risks that are exacerbated by frontier models (e.g., cyber, CBRN) and posed by nation-state actors (e.g., economic espionage by China).
- Incentivize companies to take part in this voluntary initiative by creating glide paths for them to contract with the government, including on national security projects; creating strong protections for any company information shared during these partnerships; and reducing barriers to companies' internal work related to national security domains.
- Guarantee that state-based legislation does not undermine America's innovation lead on AI. Create a sandbox for American start-ups, and provide participating companies with liability protections including preemption from state-based

⁵ Federal preemption over existing or prospective state laws will require an act of Congress.

OpenAI

regulations that focus on frontier model security (e.g., CA SB 1047). This will help keep the US public and private sectors competitive by allowing AI companies of all sizes to pursue bleeding-edge AI technology free from the regulatory uncertainty created by some state-based liability regimes.

2. Export Controls: Exporting Democratic AI

A comprehensive export control strategy should do more than restrict the flow of AI technologies to the PRC—it should ensure that America is “winning diffusion”, i.e., that as much of the world as possible is aligned to democratic values and building on democratic infrastructure. To that end, we propose that the US government consider the Total Addressable Market (TAM), i.e., the entire world less the PRC and its few allies, against the Serviceable Addressable Market (SAM), i.e., those countries who prefer to build AI on democratic rails, and help as many of the latter as possible commit, including by actually committing to deploy AI in line with democratic principles set out by the US government.

In particular, we propose maintaining the AI diffusion rule’s three-tiered framework to differentiate among countries in the global AI market, but with some key modifications that expand the number of countries in Tier I:

Tier I: Countries that commit to democratic AI principles by deploying AI systems in ways that promote more freedoms for their citizens could be considered Tier I countries.

Tier II: Limited to only those countries that have a history of failing to prevent export-controlled chips and other US-developed IP from being diverted into, or used by Tier III countries. These countries would be encouraged and supported to obtain Tier I status over time; and would be subject to more stringent security requirements in the interim.

Tier III: CCP-led China, along with a small cohort of countries aligned with the CCP, would represent its own category that is prohibited from accessing democratic AI systems.

This strategy would encourage global adoption of democratic AI principles, promoting the use of democratic AI systems while protecting US advantage. Making sure that open-sourced models are readily available to developers in these countries also will strengthen our advantage. We believe the question of whether AI should be open or

OpenAI

closed source is a false choice—we need both, and they can work in a complementary way that encourages the building of AI on American rails.

Tier I countries should include American allies, as well as those countries that are committed to democratic AI principles and that present a relatively low risk that American AI infrastructure (e.g., chips) will be diverted to non-Tier I countries. The commercial diplomacy strategy in Tier I should recognize these countries' strong history of export and customs control compliance and seek to maximally expand democratic AI systems' market share, while at the same time protecting those systems from IP theft by the PRC and other malign actors (e.g., the theft of model weights and/or chip designs, unauthorized influence or access to data center operations).

To expand market share in Tier I countries, American commercial diplomacy policy should:

- Encourage cross-border capital flows and promote software frameworks that are optimized for domestic chip design.
- Coordinate global bans on CCP-aligned AI infrastructure, including Huawei chips.
- Continue to represent American company interests in safety and security standards bodies, and encourage global regulators to adopt pro-growth safety and security policies.
- Revise the existing export control rules to eliminate country caps on compute.
- Maintain existing export license exceptions (e.g., license exception ACM) that enable exports of technology and software for technical collaboration with allies and preservation of economically critical supply chains

To protect the US-developed IP needed to operate data centers in Tier I countries, security requirements could include:

- Prohibiting relationships with Tier III countries' foreign military and intelligence services, and the use of data centers to support military/intelligence missions for Tier III nations or human rights violators.
- Banning the use of PRC-produced equipment (e.g., Huawei Ascend chips) and models that violate user privacy and create security risks such as the risk of IP theft.
- Maintaining corporate control by entities headquartered in Tier I countries.
- Implementing—and constantly modernizing—cybersecurity, model weight security, and personnel security controls that ideally are globally synced and coordinated among Tier I governments.

OpenAI

Controls on model weights—if any—should strike a balance between protecting American-developed IP and promoting the deployment of American-developed models over those developed by Tier III countries, including the PRC.

Tier II countries should include those with a history of failing to prevent export-controlled chips and other US-developed IP from being diverted into, or used by Tier III countries. Here, the commercial diplomacy strategy should still seek to expand US market share, but should do so more carefully, including by levying stronger controls on the export of US-developed AI infrastructure. At the same time, the strategy should provide transparent pathways for Tier II countries to reach Tier I status by adopting democratic AI principles and more effectively managing risks of chip diversion.

To expand American market share in Tier II countries, in addition to the steps above, the commercial diplomacy policy could be designed to leverage commercial interest in American-led AI in order to encourage investment in the US, enhanced security procedures, and encourage more countries to build on American rails:

- Establish a transparent process to evaluate countries' readiness to transition from Tier II to Tier I.
- Support countries' transition from Tier II to Tier I by helping Tier II governments strengthen their in-country security programs.
- Encourage greater economic interdependence between the US and Tier II countries.
- Incentivize public-private partnerships to rapidly mature, scale, and commercialize hardware-enabled mechanisms that could enhance in-country security controls in the future.

To protect the American-developed IP needed to operate data centers in Tier II countries, and to manage both the heightened risk of IP theft and the additional risk that export-controlled chips might be diverted from Tier II into Tier III countries, the commercial diplomacy policy also could:

- Allow the export of advanced AI chips to an end-user located in a Tier II country that meets Tier I security requirements, and that puts in place additional corporate governance controls as well as technology-enhanced protections (e.g., hardware-enabled mechanisms) against the diversion of export-controlled chips.

Tier III countries—including the PRC and any other country subject to a US arms embargo—should continue to be subject to strict export controls of AI systems, including existing export controls on advanced chips. The strategy could also expand established controls, for example, to include advanced chips that are required for large-scale

OpenAI

inference and RL training and the components used to manufacture advanced AI chips and data centers.

3. Copyright: Promoting the Freedom to Learn

American copyright law, including the longstanding fair use doctrine, protects the transformative uses of existing works, ensuring that innovators have a balanced and predictable framework for experimentation and entrepreneurship. This approach has underpinned American success through earlier phases of technological progress and is even more critical to continued American leadership on AI in the wake of recent events in the PRC. OpenAI's models are trained to not replicate works for consumption by the public. Instead, they learn from the works and extract patterns, linguistic structures, and contextual insights. This means our AI model training aligns with the core objectives of copyright and the fair use doctrine, using existing works to create something wholly new and different without eroding the commercial value of those existing works.

America has so many AI startups, attracts so much investment, and has made so many research breakthroughs largely because the fair use doctrine promotes AI development. In other markets, rigid copyright rules are repressing innovation and investment.

The European Union, for one, has created “text and data mining exceptions” with broadly applicable “opt-outs” for any rights holder—meaning access to important AI inputs is less predictable and likely to become more difficult as the EU's regulations take shape. Unpredictable availability of inputs hinders AI innovation, particularly for smaller, newer entrants with limited budgets.

The UK government is currently considering changes to its copyright regime. It has indicated that it prefers creating a data mining exception that allows rights holders to “reserve their rights,” creating the same regulatory barriers to AI development that we see in the EU.

Applying the fair use doctrine to AI is not only a matter of American competitiveness—it's a matter of national security. The rapid advances seen with the PRC's DeepSeek, among other recent developments, show that America's lead on frontier AI is far from guaranteed. Given concerted state support for critical industries and infrastructure projects, there's little doubt that the PRC's AI developers will enjoy unfettered access to data—including copyrighted data—that will improve their models. If the PRC's developers have unfettered access to data and American companies are left without fair use access, the race for AI is effectively over. America loses, as does the success of

OpenAI

democratic AI. Ultimately, access to more data from the widest possible range of sources will ensure more access to more powerful innovations that deliver even more knowledge.

We propose that the US government take steps to ensure that our copyright system continues to support American AI leadership and American economic and national security, including by:

- Shaping international policy discussions around copyright and AI, and working to prevent less innovative countries from imposing their legal regimes on American AI firms and slowing our rate of progress.
- Actively assessing the overall level of data available to American AI firms and determining whether other countries are restricting American companies' access to data and other critical inputs.
- Encouraging more access to government-held or government-supported data. This would boost AI development in any case, but would be particularly important if shifting copyright rules restrict American companies' access to training data.
- Monitoring domestic policy debates and ongoing litigation, and weighing in where fundamental, pro-innovation principles are at risk.

Generative AI models represent the next frontier of innovation, poised to revolutionize the private and public sectors, improving healthcare, education, scientific research, and so much more. If AI innovation remains protected under longstanding copyright principles, America will maintain and strengthen its role as the world leader in cutting-edge technologies and remain positioned to continue championing AI based on democratic principles with countries around the world.

4. Infrastructure: Seizing the Opportunity to Drive Growth

Today, hundreds of billions of dollars in global funds are waiting to be invested in AI infrastructure. If the US doesn't move fast to channel these resources into projects that support democratic AI ecosystems around the world, the funds will flow to projects backed and shaped by the CCP.

We propose a foundational strategy to ensure that investment in infrastructure drives economic growth that benefits all Americans; maximizes access to AI; and protects national security interests by keeping sensitive American data on American soil. This includes policies and initiatives that encourage rather than stifle developers; support a thriving AI-ready workforce and ecosystems of labs, start-ups and larger companies; and secure America's leadership on AI into the future.

OpenAI

First and foremost, building data centers is capital-intensive, particularly for newcomers seeking to compete against established hyperscalers with vast resources. We support the solutions already proposed by this Administration to ensure that sufficient capital flows to building AI infrastructure in the US:

- Investment vehicles like a Sovereign Wealth Fund.
- Government offtake and guarantees that both provide the government with the compute it needs and signal to markets that the demand will be there for American-developed AI.
- Tax credits, loans, and other vehicles the US government can direct to provide credit enhancement.

We also have proposed:

A National Transmission Highway Act, as ambitious as the 1956 National Interstate and Defense Highways Act, to expand transmission, fiber connectivity, and natural gas pipeline construction. The process for obtaining the “three Ps”—planning, permitting, and paying for approvals from federal, state, local, and tribal authorities—disadvantages America’s AI industry. Transmission lines can take 10 years or more to complete. When lines are built, parties must agree on which customers pay higher electrical bills to bear the cost of construction. In this process, delays often affect the build-out of transmission lines. Streamlining these processes and eliminating redundancies would significantly speed up infrastructure projects, keeping America’s AI sector globally competitive and securing a future of reliable, affordable energy.

Digitizing government data currently in analog form. A lot of government data is in the public domain. Making it more accessible or machine-readable could help American AI developers of all sizes, especially those working in fields where vital data is often government-held. In exchange, developers using this data could work with governments to unlock new insights that help it develop better public policies. For example, government agencies can build on the work of the US National Archives and Records Administration in using Optical Character Recognition for text searchability and AI-driven metadata tagging.

A Compact for AI among US allies and partner nations that streamlines access to capital and supply chains in ways that support AI infrastructure and a robust AI ecosystem. Participating countries would also agree to some common standards to safeguard data centers and the technology. Over time, this collaboration could expand to a global network of US allies and partners that would compete with the PRC’s AI infrastructure alliances while also strengthening security through shared standards.

OpenAI

AI Economic Zones, created by local, state and the federal government together with industry, that speed up the permitting for building AI infrastructure like new solar arrays, wind farms, and nuclear reactors. This could include creating categorical exclusions to the National Environmental Policy Act, such as a national security waiver given the global competition for AI leadership. These zones could also build on the first Trump Administration’s “Opportunity Zones” through tax incentives or credit enhancements in order to encourage private capital investment.

A nationwide AI Readiness Strategy—rooted in local communities in partnership with American companies—to help our current workforce and students become AI-ready, bolster the economy, and secure America’s continued leadership on innovation. Maintaining American leadership in AI means ensuring we have an experienced, trained professional workforce working across the AI supply chain, including construction worker management, HVAC technicians, and electricians. Government should ensure this training is accessible and affordable, such as by:

- At the federal level, expanding 529 savings plans to cover more AI supply chain-related training programs—including for construction, HVAC technicians, electricians, as well as AI researchers and developers—by amending Section 529 of the Internal Revenue Code or broadening the SECURE Act’s provisions.
- At the federal or state level, incentivizing AI supply chain companies to work with a backbone organization to understand the workforce needs of AI supply chain companies, develop a pipeline of training programs that help companies meet those needs, and coordinate with labor unions, community colleges, and trade associations to build and operate that training pipeline.

Creation of AI research labs and workforces aligned with key local industries by requiring AI companies to provide meaningful amounts of compute to public universities to equitably scale the training of a homegrown AI-skilled workforce. For example, one state could establish a hub dedicated to applying AI in agriculture while another develops centers focused on integrating AI into power production and grid resilience.

Using the Defense Production Act (DPA) Title I to manage supply chain risk by designating gas turbines, rankine cycle turbines, high-voltage transformers, or switchgear for data centers as “rated orders.” This prioritization could significantly shorten timelines for data center power infrastructure projects.

OpenAI

5. Government Adoption of AI: Leading by Example

AI adoption in federal departments and agencies remains unacceptably low, with federal employees, and especially national security sector employees, largely unable to harness the benefits of the technology.

The government should encourage public-private partnerships to enhance government AI adoption by removing known blockers to the adoption of AI tools, including outdated and lengthy accreditation processes, restrictive testing authorities, and inflexible procurement pathways. Specifically, we recommend:

- *Modernizing cyber security rules for cloud-based applications.* The government's current processes for AI providers to comply with federal security regulations—primarily through the Federal Risk and Authorization Management Program (FedRAMP)—takes 12 to 18 months, compared to the one- to three-month commercial standard, with no clear evidence of additional protection for government data. The government should modernize FedRAMP by establishing a faster, criteria-based path for approval of AI tools. Criteria could include Foreign Ownership, Control, or Influence (FOCI) approval; Facilities Clearance (FCL) status; US incorporation; a first-party AI model that ranks in the top 20 of a recognized evaluation framework (for example, MMLU, or Massive Multitask Language Understanding); SOC 2 (System and Organization Controls 2) accreditation; and a recent third-party penetration test with all findings addressed.
- *Accelerating AI testing and experimentation.* The government should allow federal agencies to test and experiment with real data using commercial-standard practices—such as SOC 2 or International Organization for Standardization (ISO) audit reports—and potentially grant a temporary waiver for FedRAMP. AI vendors would still be required to meet FedRAMP continuous monitoring requirements while awaiting full accreditation. Combined with standard due diligence before actual use, this approach could allow agencies to access new AI services roughly 12 months earlier while maintaining compliance with federal security requirements.
- *Enabling rapid procurement mechanisms.* Once new security and testing approaches are in place, agencies must also have quicker, more direct routes to procure and deploy frontier AI tools. The government should continue to evaluate Other Transactional Authorities (OTAs), Commercial Service Offerings (CSO) or other procurement paths in order to access technology from frontier AI labs, not just their legacy IT providers. We are encouraged by the Department of Defense's recent efforts to [Modernize Software Acquisition](#).

OpenAI

Enabling federal agencies to quickly acquire consumer-focused models is not enough, however. The government also needs to pursue and fund bespoke national security pilot projects for which there may be no commercial market by:

- *Partnering with industry to develop custom models for national security.* The government needs models trained on classified datasets that are fine-tuned to be exceptional at national security tasks for which there is no commercial market—such as geospatial intelligence or classified nuclear tasks. This will likely require on-premises deployment of model weights and access to significant compute, given the security requirements of many national security agencies.
- *Acting now to fund these projects and secure this compute*—enabling industry partners to secure chips, transformers, and begin construction, and ensuring that this compute comes online at the pace that innovation and geopolitical competition require.

Lastly, frontier AI labs need Facility Clearances (FCL) to work directly with the national security enterprise on these pilot projects and custom models. The government should:

- *Expedite FCL for frontier AI labs committed to supporting national security.* The process for obtaining a FCL can take a year or longer. Given the rapid pace of AI development, the government should start prioritizing deeper collaboration with frontier AI labs as soon as possible.

We look forward to discussing the above proposals with the Office of Science and Technology Policy as we continue to build on our relationship with the US government and work toward AI that benefits everyone.

About OpenAI

OpenAI's mission is to ensure that as AI advances, it benefits everyone. We're building AI to help people solve hard problems because by helping with the hard problems, AI can benefit the most people possible—through more scientific discoveries, better healthcare and education, and improved productivity. We're off to a strong start, creating freely available intelligence being used by more than 400 million people around the world, including 3 million developers. We believe AI will scale human ingenuity and drive unprecedented economic growth and new freedoms that help people accomplish what we can't even imagine today.