



October 27, 2025

TO: Michael Kratsios
Executive Director
Office of Science and Technology Policy
Eisenhower Executive Office Building
1650 Pennsylvania Avenue
Washington, DC 20504

FROM: Christopher Lehane
Chief Global Affairs Officer
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 AI policies and associated documents without attribution.

The United States currently leads the world in developing artificial intelligence—an advantage that ensures that AI will be built to benefit as many people as possible, not just a few.

A year ago, we described how the national security imperative to lead the world in AI also presents a once-in-a-century opportunity to strengthen our economy and modernize America's industrial base. Seizing this opportunity will catalyze domestic manufacturing, update our energy grid, create more well-paid jobs, and generate more revenue—while also strengthening American leadership in the world.

According to an OpenAI internal analysis, just the first \$1 trillion invested in AI infrastructure could result in more than 5% in additional GDP growth over a 3-year period.

An analysis commissioned by OpenAI of our own plans to build AI infrastructure in the US—with six Stargate sites already underway in Texas, New Mexico, Ohio, and Wisconsin, and more to come—finds that our plans over the next five years will require an estimated 20% of the existing workforces in skilled trades such as specialized electricians and mechanics. The country will need many more electricians, mechanics, metal and ironworkers, carpenters, plumbers and other construction trade workers than we currently have. Americans will have new opportunities to train into these jobs and gain valuable, portable expertise.

But limits on how much electricity the US can generate to power AI development threaten both our ability to seize this once-in-a-century opportunity, and our advantage on the most consequential technology since electricity itself. In pursuit of its goal to overtake the US and lead the world on AI by 2030, the People's Republic of China (PRC) has built real momentum in energy production—treating capacity as the foundation of industrial competitiveness. In 2024, the PRC added 429 GW of new power capacity, more than one-third of the entire US grid, and more than half of all global electricity growth. The US contributed just 51 GW, or 12%.

OpenAI

For the US, unlocking electrons will unlock our greatest national economic opportunity since electricity drove the latter half of the Industrial Age. Electricity is a strategic asset—it underpins the US AI advantage and in turn, our economic competitiveness, technological leadership, and national security. And AI, as the foundational technology of our time, is itself becoming a powerful tool for innovating with energy. It's already improving how we model the grid, forecast demand, and design next-generation systems. Building more capacity to power AI will in turn make our energy generation more efficient.

OpenAI was pleased to see the Trump Administration's AI Action Plan recognize how critical AI is to America's national interests, and how maintaining the country's lead in AI depends on harnessing our national resources: chips, data, talent and energy. We also were pleased to see that last week, the Department of Energy moved to streamline state-by-state processes that have become one of the biggest barriers to building the energy infrastructure required for AI. We want to recognize the work of the Office of Science and Technology Policy (OSTP) to engage stakeholders across AI to inform the Administration's policymaking.

As we look ahead to 2026 and beyond, we recommend that OSTP prioritize closing the "electron gap" between the US and PRC by setting an ambitious national target of building 100 GW a year of new energy capacity. This amount would both account for baseline demands and grid retirements, and fuel American AI dominance, ensuring that US companies have the power they need and energy becomes cheaper and more reliable for communities across the country.

The US has a history of meeting the moment by thinking big, acting big, and building big. Communities everywhere, not just in Silicon Valley, will benefit from the economic upsides of the Intelligence Age—in new jobs and revenue from construction, and in the long-term build-out of chips fabrication facilities and the manufacturing of turbines, copper wire, transformers and other essential components.

The vastness of the opportunity creates room for the federal government to both step back and let markets work, and also lean in. It can step back by streamlining and modernizing outdated and onerous regulations to unlock energy innovation while continuing to protect our national security, including through its own adoption of AI. It can lean in by investing in our next great era of manufacturing and in the workforce to run it, especially in rural and underserved communities. Specifically, to achieve 100 GW in new energy per year and the benefits that come with it, the federal government should:

1. strengthen America's industrial base by investing in manufacturing across the country;
2. unlock more energy by removing or modernizing regulations;
3. equip American workers and their families by helping to fund workforce development at the state and local level, and scaling AI education; and,
4. at the same time, continue to ensure that frontier AI systems protect American national security interests, including through federal agency adoption.

OpenAI

OpenAI is committed to doing our part. Through our Stargate initiative, the six sites we've already announced bring Stargate to nearly seven GW of planned capacity and over \$400 billion in investment over the next three years. This puts us on a clear path to securing the full \$500 billion, 10 GW commitment we announced in January by the end of 2025—ahead of schedule. We continue to evaluate additional sites across the US.

In 2026 and beyond, we'll build on that progress by strengthening the broader domestic supply chain—working with US suppliers and manufacturers to invest in the country's onshore production of critical components for these data centers. We will also develop additional strategic partnerships and investments in American manufacturing to specifically advance our work in AI robotics and devices.

We are committed to building new pipelines of workers by offering certifications for different levels of AI fluency, while connecting our Stargate sites with local workforce and education partnerships through our Certifications and Jobs Platform, starting in 2026 with our site in Abilene, TX. The country will need many more specialized workers than we have now, creating opportunities for Americans to train into these jobs and gain valuable, portable expertise.

We're also committed to building AI infrastructure that helps modernize the grid and generate more energy for everyone. In many places, we're not just plugging into the grid; we're building new energy supply, including natural gas and solar projects behind the meter, as well as new grid-connected capacity like the Wisconsin Stargate site we just announced with Oracle. Below, we detail ways in which we're pursuing our goal to become a net contributor to the regions we operate in—adding capacity, improving resilience, and helping keep energy reliable and affordable for everyone.

OpenAI's mission is to ensure that artificial general intelligence benefits all of humanity. We're humbled that since our March 2025 response to OSTP to help inform its AI policy agenda, our weekly user base has doubled from over 400 million to over 800 million, with most users accessing our tools for free. We see this reindustrialization as a foundational way for the US to “predistribute” the economic benefits of the Intelligence Age from the very start. As with the wheel, the printing press, the combustion engine, electricity, the transistor, and the internet, if we make it possible for the US to occupy the center of this Age, it will lift all Americans regardless of where they live, not just cohorts in certain parts of the country.

We remain committed to working with OSTP and others across the federal government toward advancing the enclosed proposals, and we look forward to discussing them with you.



Chris Lehane
Chief Global Affairs Officer

OpenAI

1. *Strengthen America's Industrial Base*

The rapid expansion of AI has exposed vulnerabilities in the US industrial base. Every stage of the AI buildout—from critical minerals and semiconductor fabrication to grid components and data center construction—relies on globally concentrated inputs or constrained manufacturing capacity. Bottlenecks across this chain have already delayed projects, threatening America's ability to execute quickly enough to maintain global leadership. The Cardinal-Hickory Creek transmission line in the Midwest took 13 years to complete; interconnection queues have faced significant delays.

With domestic production unable to meet surging demand from AI data centers, lead times for critical equipment now often exceed two years, threatening both compute expansion and grid reliability. Upstream, these systems depend on critical minerals like copper, aluminum, electrical steel, rare earths, and semiconductors—which remain dominated by a handful of foreign suppliers, exposing the US to geopolitical and supply risks.

Yet, this challenge presents a historic opportunity to reindustrialize the country. The physical investment required to secure AI supply chains can anchor a new generation of American manufacturing. This is an area where OSTP and the federal government should lean in, through the following proposals.

Expand tax credits that accelerate American manufacturing. The Administration has already taken critical steps to strengthen American manufacturing by extending the Advanced Manufacturing Investment Credit (AMIC) for semiconductor fabrication. OSTP should now double down on this approach and work with Congress to further extend eligibility to the semiconductor manufacturing supply chain; grid components like transformers and specialized steel for their production; AI server production; and AI data centers. Broadening coverage of the AMIC will lower the effective cost of capital, de-risk early investment, and unlock private capital to help alleviate bottlenecks and accelerate the AI build in the US.

Counter the PRC by de-risking US manufacturing expansion. To provide manufacturers with the certainty and capital they need to scale production quickly, the federal government should also deploy grants, cost-sharing agreements, loans, or loan guarantees to expand industrial base capacity and resilience. This financial support would help counter the PRC, especially in instances where it is distorting the market, such as in copper, aluminum, electrical steel, rare earth elements, and semiconductor inputs. Direct funding could also help shorten lead times for critical grid components—transformers, HVDC converters, switchgear, and cables—from years to months. Initial investments could be made using existing authorities such as the Defense Production Act Title III and the Department of Energy's Loan Programs Office.

Accelerate transmission line construction. High upfront costs, long permitting timelines, and fragmented regional investment have slowed the expansion of high-voltage transmission needed to connect new generation and growing AI-driven loads. The Administration should

OpenAI

accelerate transmission line construction by using existing tools—such as DOE’s Loan Programs Office, the US Department of Agriculture’s Rural Utilities Service loans, DOE’s Grid Resilience and Innovation Partnerships (GRIP) program grants, and Power Marketing Administration financing—to mobilize capital and speed project delivery. The Federal Energy Regulatory Commission (FERC) should further strengthen coordination and cost-sharing for interregional lines through rulemaking under its existing authority, and work with Congress to expand its backstop-siting authority to overcome state-level permitting barriers. Together, these actions would ensure fair cost distribution, reduce delays, and enable a faster, more reliable grid buildout.

Create a Strategic Reserve for the raw materials essential to AI infrastructure. Drawing on precedents like the Strategic Petroleum Reserve, Congress and the Administration should establish a strategic reserve of copper, aluminum, processed rare earth elements, and raw semiconductor materials. This reserve would build off the important steps that the Administration has already taken to invest in critical minerals (e.g., through its investment in MP Materials and efforts to establish a \$1 billion critical minerals reserve). Applying the same concept to the materials that are essential to AI infrastructure would provide domestic producers with the certainty they need to expand production. It would also reduce our dependence on the PRC, stabilize prices during shortages, and ensure secure access to materials vital for US AI infrastructure. Initial steps could be taken using existing authorities such as the DPA Title III, the Strategic and Critical Materials Stock Piling Act, and DOE’s Loan Programs Office.

2. Unlock More Energy By Modernizing Regulations

The US cannot maintain its AI advantage if the “electron gap” with the PRC persists and our ability to build new energy capacity is trapped in a regulatory chokehold. Although interconnection queue delays have slowed or stalled projects in the past, this is an area where the federal government can step up by stepping back. President Trump’s Executive Order 13807 (“One Federal Decision”) set appropriately ambitious 2-year targets for environmental reviews and sped up permitting using emergency authorities. To continue unlocking American energy, we encourage OSTP to direct additional steps to modernize regulations.

Bring new capacity onto the electric grid faster. For example, the government can tighten FERC Order Number 2023 timelines by instituting automatic enforcement of late submissions and standardizing timelines for extensions. The government also can create a fast track for shovel-ready projects (e.g., storage, repowerings) to allow simple AI infrastructure projects to skip the line.

Accelerate transmission build out and permitting. The federal government should prioritize projects that connect high-generation regions to major load centers. Using existing tools—such as the DOE Coordinated Interagency Transmission Authorization and Permits (CITAP) process, National Interest Electric Transmission Corridor (NIETC) designations, FERC’s backstop-siting

OpenAI

authority, and FAST-41 permitting coordination—federal agencies can move critical projects from planning to construction faster. Combined with GRIP grants that support advanced conductors and grid-enhancing technologies, these actions would cut project timelines and strengthen the national transmission network.

Expand use of curtailable load resources and modernize interconnection policy. To strengthen grid reliability and expand capacity for AI and other flexible loads, FERC should allow more demand-side participation in wholesale markets and speed up interconnection for large loads that can curtail. We welcomed the news last week that DOE recommended to FERC that it assert jurisdiction and create standardized rules for large load interconnections. Indeed, such rules could build on the Electric Reliability Council of Texas’s successful “Connect and Manage” model, transmission operators could allow new projects to connect before full transmission upgrades are complete, provided that they can reduce consumption or output (curtail) before power demand approaches the limits of the grid. This approach would increase use of existing grid capacity, reduce interconnection wait times, and reduce costs for ratepayers while maintaining reliability across the Bulk Electric System.

Accelerate permitting under the Clean Water Act (CWA) and Clean Air Act (CAA). The government should shorten timelines for CWA §401 certifications, expand and standardize nationwide permits for linear energy work, and create a standardized permitting process for data center campuses that use advanced cooling processes. Recent announcements about the intent to revise the New Source Performance (NSP) provisions on starting actual construction under the CAA represent a positive first step that could be furthered by stabilizing PM2.5 implementation, creating a fast track New Source Review (NSR) playbook for reliability-critical projects, and providing uniform rules for backup generators at data centers.

Leverage AI to speed up federal permitting and environmental review processes. OpenAI is already working with the US National Labs to deploy AI to improve the accuracy and speed of federal permitting. OSTP could build upon this early progress by directing AI-enabled review of process bottlenecks, especially under the Clean Water Act, Clean Air Act, and Endangered Species Act. It could also direct agencies to have existing authorities launch AI-assisted permitting and environmental reviews. OSTP could also encourage the Council on Environmental Quality to work with AI companies to speed the permitting process.

Speed up the permitting and environmental review processes for projects that build on federal lands or use federal funds. We have endorsed bipartisan Congressional proposals like the SPEED Act that would address longstanding issues with NEPA delays. Where AI infrastructure construction occurs on federal lands or using federal funds, the federal government should accelerate the National Environmental Policy Act (NEPA) review process. For example, it can seek sped-up or alternative arrangements for NEPA by considering the national security implications of AI infrastructure construction. It also should create categorical exclusions to speed up reviews when construction would not have a significant impact on the environment (e.g., DOE recently added a categorical exclusion for micro-reactors on federal sites).

OpenAI

At OpenAI, we are bringing new power online to support sites like our flagship Stargate campus in Abilene, and in grid-connected regions we're working with local utilities to expand capacity and strengthen stability. Our data centers are designed to be curtailable—reducing their draw or even returning power during peak demand, helping to protect reliability and avoid higher costs for consumers. Across all our US Stargate locations, our operations focus on optimizing both energy and water. And in areas where we build behind-the meter power, we are pushing to connect to the grid so that surplus energy can support local communities. Our goal is to become a net contributor to the regions we operate in—adding capacity, improving resilience, and helping keep energy reliable and affordable for everyone.

3. Equip American Workers and Their Families for the AI Economy

OpenAI supports the workforce policies proposed in the Trump Administration's AI Action Plan. We recently released our [Workforce Blueprint](#), a comprehensive set of proposals to build on that foundation with steps that can be taken today, and steps OpenAI already is taking, to make AI and the ability to use it available to as many people as possible:

- Most of our over 800 million weekly users are using our AI for free.
- Nearly 2.5 million Americans already have engaged with our AI literacy content through the OpenAI Academy learning platform.
- Through our recently announced [OpenAI Certifications](#) program, we'll help 10 million Americans improve their AI skills by 2030 and give employers the confidence they need to hire with trust. Launch partners ranging from Walmart to the Texas Association of Business, representing thousands of smaller businesses across the state, will help ensure that enterprises of all sizes are able to provide training for their employees on OpenAI's best-in-class tools.
- Through the [OpenAI Jobs Platform](#), launching in 2026, we'll help future-proof the workforce by giving millions of people a path to stability and growth through better jobs and long-term career security.

To close the “electron gap” specifically, the US needs not only steel and concrete, but also a skilled workforce at a scale not seen in generations. This includes mechanical, electrical, and plumbing trades, and engineering and management specialists across the infrastructure stack. Reaching that scale means significantly increasing our current training pipeline for core electrical and construction trades, and quickly expanding specialized training in areas like transformer manufacturing, HVDC systems, high-voltage cable installation, and grid commissioning.

Our Workforce Blueprint suggests an approach for building these training pipelines by connecting AI infrastructure companies with community colleges and technical and trade schools through local “AI Hubs.” These hubs should be managed at the state and local level but should leverage federal funding (e.g., the Workforce Innovation and Opportunity Act) and use the workforce intermediary strategy endorsed by the AI Action Plan. Workforce intermediaries connect

OpenAI

employers to local training providers such as community colleges and trade schools, and OpenAI is proud to connect our Stargate data center campuses with community workforce and education partnerships through our Certifications program and Jobs Platform, starting in 2026 with our site in Abilene, TX.

We also envision hubs serving as structured spaces for training and peer learning, including among parents. In addition to offering physical spaces for in-person and hands-on education, hubs can also host regular “AI show-and-tell” events where workers, employers, and training providers can share their experiences using AI—what worked, what failed, and what they learned. Similarly, hubs can offer a place for parent workshops that could teach parents how to use parental controls and monitor safe AI use; demonstrate practical ways AI can save time, improve decision-making, or foster creativity; and offer space for open Q&A and discussion to demystify the technology. Ideally communities would participate in the co-design of these sessions, to provide the expertise on how best to tailor training for their residents.

4. Ensure That Frontier AI Systems Advance US National Security Interests

The federal government is uniquely positioned to ensure that frontier AI systems protect American national security interests, including through federal agency adoption. National security is also an area where the federal government should both lean in and enable progress—by creating opportunities for public-private partnerships that advance US security and competitiveness, and by removing unnecessary regulatory barriers to safety testing and government adoption.

We also believe that on matters of national security and global competitiveness, the federal government should play a singular leadership role. A patchwork of state-level national security regulations risks bogging down innovation and weakening America’s global leadership position. At the same time, states can serve as valuable “laboratories for democracy” by piloting complementary policies in areas such as child and teen safety, AI literacy and education, and workforce development. The following proposals suggest opportunities for the federal government to lean in, step back, and ensure that state-based regulation doesn’t harm US national security interests.

Remove regulatory roadblocks to safety testing. Even at the federal level, regulations can impede the safe development and deployment of AI systems. To clear these roadblocks, OSTP should direct relevant departments to provide exemptions, authorizations, or clarifications to allow AI companies to responsibly and efficiently evaluate and improve system safety. For example:

- The Bureau of Industry and Security (BIS) and the Directorate of Defense Trade Controls (DDTC) should provide general authorizations, or confirm that such authorizations would not be required, for “exports” and “deemed exports” related to AI safety work in chemical, biological, radiological, nuclear (CBRN) and other sensitive

OpenAI

domains. This would ensure that US companies can conduct essential testing and other safety work with vetted personnel without having to go through a lengthy license approval process.

- The Department of Justice should provide responsible AI companies with immunity in instances where they are conducting legitimate child safety red-teaming and safety evaluations under strict containment and using protocols that are cleared by DOJ's Child Exploitation and Obscenity Section or similarly reputable entities such as the National Center for Missing and Exploited Children (NCMEC).

Incentivize companies to partner with the Center for AI Standards and Innovation (CAISI).

National standards for frontier AI models are still nascent, and federal regulations should not outpace consensus-based standards. To speed up development of standards, the federal government should encourage more companies to voluntarily partner with the CAISI, including by working with Congress to provide participating companies with liability protections such as preemption of state laws and regulations that focus on frontier model security and safety if companies partner with the CAISI. This will help keep the US public and private sectors competitive by allowing AI companies of all sizes to pursue cutting-edge AI technology while also ensuring that they continue critical safety testing—all without being hampered by the regulatory uncertainty created by some state-based liability regimes.

Equip the CAISI to serve as the single, efficient “front door” to the federal government. The CAISI should continue to centralize technical expertise, streamline risk assessment and best practices development, and reinforce US competitiveness and national security. By positioning the CAISI as the hub for evaluation, red-teaming, intelligence integration, and interagency coordination, OSTP can efficiently support US companies while mitigating the risks posed by frontier AI systems. In particular, we recommend:

- Fully resourcing the CAISI with personnel and technical capabilities—including experts on frontier AI models—to strengthen its ability to serve as a lasting federal center of excellence on AI standards and innovation.
- Institutionalizing CAISI-led comparisons of foreign AI systems and publishing interagency advisories for procurement and critical infrastructure operators.
- Integrating unique CAISI technical expertise of AI systems into US government national security assessments of nation state and other sophisticated malicious uses of AI.
- Reporting annually on CAISI's evaluation, red-teaming, and standards development activities, including voluntary industry partnerships.
- Coordinating with external industry-supported groups such as the Frontier Model Forum to facilitate sharing technical information and best practices to facilitate development of standards and shared resources.
- Exempting sensitive technical data voluntarily shared with the CAISI and partner agencies from Freedom of Information Act (FOIA) disclosure, in order to encourage robust public–private collaboration while preserving transparency where it matters.
- Directing the CAISI to develop a strategy for working with start-ups to share useful information to advance national security while enabling their innovation to flourish.

OpenAI

- Encouraging the CAISI to guard against burdensome international regulations by coordinating with the National Institute of Standards and Technology (NIST) Information Technology Laboratory (ITL), the State Department, and other agencies to ensure consistent US representation in key international standards bodies, while issuing clear guidance to US companies on how to engage effectively.

Ensure that any new federal regulations are future-proofed. Make sure that any future federal AI regulation is flexible and performance-based (i.e., focused on objectives, not over-specifying designs), with mechanisms that: (1) reference evolving, consensus standards; (2) provide safe harbors for compliance with recognized standards; (3) use accredited conformity assessment rather than prescriptive checklists; (4) enable pilots and regulatory sandboxes to validate methods before codification; and (5) include periodic review and sunset clauses so requirements keep pace with science.

Continue improving and expanding FedRAMP 20X. The Trump Administration has made real progress in speeding up how new technologies get approved for government use. FedRAMP's new "20X" process and AI prioritization effort have replaced outdated paperwork-heavy reviews with a faster, more flexible system. Instead of static, one-time audits, the 20X process focuses on real security outcomes and uses automation to keep systems safe in real time. The AI prioritization track also gives companies direct support and quick feedback during reviews, helping them fix issues early and get approved faster. These changes have made FedRAMP much more efficient, and the Administration should continue to expand and strengthen these improvements.

Streamline the Department of War's security review process. OSTP should direct the Department of War (DOW) to modernize its security approval process for new technologies—governed by the Security Requirements Guide (SRG) and Cybersecurity Maturity Model Certification (CMMC)—which is still too slow and paperwork-heavy. The current system was built for older software and can delay AI projects for months, even when they pose little or no security risk. By moving toward a faster, more flexible system like FedRAMP's 20X process, the Department could use automation and real-time checks instead of static reviews. A streamlined framework could allow temporary approvals for low-risk AI projects, recognize security reviews done by other agencies, and speed up testing while maintaining strong cybersecurity protections. This would help the Department adopt and learn from new AI technologies more quickly without compromising security.

Simplify and align DOW cybersecurity rules. Today, companies face overlapping and confusing rules and processes under FedRAMP 20x, SRG, CMMC, NIST 800-53, and NIST 800-171, leading to duplicated work and uncertainty about which standards and processes individual DOW offices will accept. DOW should create a single, unified framework for cybersecurity compliance that aligns requirements and processes, reduces red tape and complexity, and allows companies to focus on innovation rather than compliance.

OpenAI

Modernize accessibility rules for AI tools. Federal accessibility rules under Section 508 ensure that technology is usable by everyone—a vital goal. Because AI systems update so often, OSTP should direct an update to these rules that would create a flexible, tiered process—one that allows quick pilots for low-risk projects while working with experts and users to build lasting accessibility standards. This would uphold the spirit of Section 508—equal access—while allowing AI innovation to move forward responsibly.

Align the NIST AI Risk Framework with existing laws. OSTP should direct the National Institute of Standards and Technology (NIST) to publish clear “crosswalks” showing how the updated AI Risk Management Framework connects to existing laws like HIPAA, Reg SCI, and other safety standards. This would help organizations see how using the NIST framework can also meet legal requirements. This also would benefit smaller organizations building AI tools. NIST should also create a model “presumption of adequacy,” so regulators can treat companies that follow the framework as already meeting core compliance standards—making the process more predictable and efficient.

Build the secure backbone required for enterprise-level national security AI. The government faces a growing deficit in GPU capacity across its classified cloud environments, constraining the nation’s ability to deploy AI at speed and scale for critical national security missions. OpenAI proposes a “Classified Stargate” initiative to help meet this need—mobilizing private capital alongside government partners to establish accredited, classified data centers purpose-built for government AI. This model would enable the government to operate outside of hyperscaler constraints while providing the trusted, high-performance infrastructure essential for the next generation of defense and intelligence capabilities.

Update data residency and sovereignty requirements. Federal data residency expectations require data to remain within US boundaries. But inconsistent interpretations across agencies impedes the use of commercial AI infrastructure for testing and evaluation. For low-risk or non-sensitive datasets, these restrictions also slow experimentation without meaningfully improving security. OSTP should consider updating these requirements and enabling temporary or conditional exceptions for approved pilot projects with clear criteria for what qualifies as non-sensitive data. This approach preserves national data protections while enabling faster, secure experimentation and informed policymaking on long-term data sovereignty standards.

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

About OpenAI

Artificial intelligence is an innovation like electricity—it will change how we live, how we work, and how we engage with one another. OpenAI’s mission is to ensure that artificial general intelligence benefits all of humanity. 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 800 million people around the world, including 4 million developers. We believe AI will scale human ingenuity and drive unprecedented productivity, economic growth, and new freedoms that help people accomplish what we can’t even imagine today.