OpenAI

Via Electronic Submission
National Telecommunications and Information Administration (NTIA)
U.S. Department of Commerce
1401 Constitution Avenue NW
Washington, DC 20230

Subject: Response to NTIA Request for Comments on Bolstering Data Center Growth, Resilience, and Security (Docket Number 240823-0225)

OpenAI welcomes the opportunity to respond to the National Telecommunications and Information Administration's Request for Comments on "Bolstering Data Center Growth, Resilience, and Security."

The United States leads the world today in the development of artificial intelligence because of decisions decades ago to install fiber-optic cables, coaxial lines and other broadband infrastructure that put the country at the forefront of the early digital revolution. The 1996 Telecommunications Act, with bipartisan support from forward-thinking lawmakers, reinforced infrastructure as a national priority, and the country's resulting leadership in AI demonstrated that infrastructure is destiny.

As revolutionary as electricity, and promising similarly distributed access and benefits, AI can power a reindustrialization across the United States, extend America's global competitiveness, and boost national, state and household finances for the long term. Investment to extend the U.S. lead in AI can yield tens of thousands of jobs; significant growth in GDP; a modernized energy grid and cleaner sources of energy like nuclear power; and a state-of-the art network of semiconductor manufacturing facilities – invigorating local economies across the country.

With an estimated $175 billion in global infrastructure funds waiting to be committed, the question is not whether that funding will flow, but where. If it doesn't flow into U.S.-backed global infrastructure projects that advance a global AI that spreads the technology's benefits to the most people possible, then it will flow to China-backed projects that leverage AI to cement and expand autocratic power. There is no third option.

OpenAI recently engaged outside experts to forecast potential job gains and GDP growth that would result from building 5GW data centers in a sampling of states. We found that constructing and operating a single 5GW data center could create or support about 40,000 jobs – in construction and maintenance, restaurants and retail, and other industries that would serve the new workers — and contribute between $17 billion and $20 billion to a state's GDP.

These numbers highlight the importance of getting data center policy right. To that end, we provide the following feedback to help policymakers invest in AI infrastructure growth, resilience and security.

OpenAI

Current and Future Challenges and Opportunities for Data Centers in the U.S. [Question 1]

- **Challenges:**
    - **Power Supply and Infrastructure:** Data centers face challenges in accessing reliable, affordable power and power grids require substantial upgrades. Even if overall grid capacity grows by 2% annually, utilities still need to add well in excess of 100 GW of peak capacity to a system that currently handles around 800 GW at peak capacity. The increase in power demand will also likely be hyper-localized, potentially requiring a doubling of grid capacity over the next decade in areas with a concentration of data centers. Approximately $50 billion of investment through 2030, or roughly $7 billion annually, is needed to facilitate the new power generation alone. North America as a whole is expected to add ~200 GW to the grid by 2030. The economically transformative AI models that benefit the United States and the rest of the world will use tens of GW of this additional power in coming years and may require 100 GW of that power by 2030. Between generation, transmission, and distribution needs, utility companies may spend nearly 40% more from 2024-2027 relative to the prior four-year period, amounting to roughly $140 billion on average annually.[1] The implementation of the Inflation Reduction Act provides tailwinds by decreasing the need for long-haul transportation gear and distribution cables, but implementation is far from complete.
    - **Regulatory Hurdles:** The process for obtaining the “three Ps” – planning, permitting, paying for approvals from federal, state, local, and tribal authorities – can be lengthy and complex, particularly for energy infrastructure. Natural gas and solar plants can be built in under two years, as can interconnects, but transmission lines can take 10 years or more to complete. When transmission lines are built, parties must agree on cost allocation — deciding which customers should pay higher electrical bills to bear the cost of construction. Disagreement over cost allocation of long-distance transmission lines can delay these lines and increase times in the “interconnection queue.” These challenges can be compounded by delays caused by negotiations associated with hot-button political issues such as local electricity rates, power allocation, and carbon offsets.
    - **Environmental Concerns:** The energy-intensive nature of data centers incentivizes increased investment in carbon capture and storage, but this will take time to scale. Companies will need to manage their AI data center

[1] Goldman Sachs Global Investment Research, Issue 129

infrastructure, and related power footprint expansion, with environmental and offset commitments.
    - **Supply Chain Vulnerabilities:** Data centers rely on a global supply chain for critical components such as transformers, cooling systems, servers, and networking equipment. Memory is also an increasingly critical input at scale, especially HBM (high-bandwidth memory) for AI superclusters. This is increasingly a supply constrained input with overseas sources, on par with AI GPUs over the coming years. Disruptions or shortages in advanced packaging, as well as fiber optics, can hinder expansion. Packaging tech concentration of CoWoS (Chip on Wafer on Substrate), and other technologies in Taiwan, for example, is an additional vulnerability beyond the latest nodes of semiconductor fab infrastructure.
    - **Workforce Shortages:** There is a shortage of skilled workers, such as network engineers, cybersecurity professionals, and construction labor, which could slow down data center construction and operation.
- **Opportunities:**
    - **Global Partnerships**: America's leading AI companies are building partnerships around the globe that effectuate cost and regulatory arbitrage in a manner that considers the competitive advantages of each country when evaluating technology investments. The United States has a unique advantage because of American leadership in entrepreneur-driven AI software, hardware, and public/private research innovation. This advantage can be leveraged more through proactive partnerships driven by differentiated market realities of the U.S. entrepreneurial/government partnership dynamic versus these relationships in the EU and China.
    - **AI and Emerging Technologies:** New technologies will produce increased efficiencies. These include: AI processor design, advanced packaging combined with high-bandwidth memory, on chip liquid cooling, optical networking, and LLM load management. The massive scale of this multi-geography 'AI data center fabric' will result in cost advantages to both the operators and their customers that will provide efficiency gains over time. Government has a role to play in accelerating economies of scale in these technologies by promoting industrial and technology clusters through financial and regulatory incentives.
    - **Modernization and Efficiency Gains:** Data centers can capitalize on opportunities to modernize infrastructure by implementing energy-efficient technologies (e.g., advanced cooling, renewable energy integration) to reduce operational costs and environmental impact. The process of upgrading the half trillion dollar data center industry with AI upgrades will likely require more than $2 trillion in investments over the coming decade, according to Nvidia.

- **Increased Government Support:** Federal and state governments may offer incentives, grants, or regulatory support to encourage data center expansion, especially if aligned with clean energy or digital infrastructure goals. It is a particular strength of the United States that there is bipartisan government engagement and interest in accelerating regulatory and financial support at both the federal and local levels.

Critical Market Considerations for Data Center Modernization or Expansion [Question 2A]

- **Telecommunications Infrastructure:** Access to high-speed broadband and robust telecommunications infrastructure is essential to support large-scale data processing and low-latency communication. In particular, 'dark fiber' availability in large quantities over land and sea is a critical variable in constructing the necessary 'AI data center' fabric. Also important is their location in proximity to the power sources to minimize regulatory and deployment delays from additional grid extensions.
- **Access to Skilled Workforce:** Proximity to areas with a strong pool of technical talent, such as network engineers and data scientists, is important for staffing and operational efficiency. Educational institutions across the United States can provide training and certification of new categories of 'data center personnel', which will be needed in large quantities. These new categories of workers blend traditional blue- and white-collar work categories and provide opportunities for the government to catalyze new types of jobs.
- **Government Incentives:** Access to government incentives, such as grants and local financial incentives, can significantly reduce costs and accelerate timelines.
- **Land and Water Availability:** Sufficient land and water resources are necessary for building and cooling large data centers, particularly in regions with less developed infrastructure.
- **Power Grid Connectivity and Costs:** Connectivity to a reliable power grid with competitive electricity rates is crucial, especially in regions with high energy costs. AI data centers adjacent to recommissioned and extended nuclear plants in the United States may offer connectivity cost and time advantages compared to traditional customers requiring meaningful investments in transmission infrastructure.

Role of Competition Among Hyperscalers and Barriers for New Entrants [Question 2B]

- **Barriers for New Entrants:**
  - **High Capital Requirements:** Building data centers is capital-intensive, particularly when competing against established hyperscalers with vast resources.

    - **Regulatory and Compliance Barriers:** New entrants face regulatory hurdles and must comply with a myriad of local, state, and federal regulations, which can be costly and time-consuming.
    - **Access to Key Resources:** Established players have long-standing relationships with suppliers and access to critical resources, such as fiber optic networks, power, water and land.
    - **Customer Switching Barriers:** Customers may face challenges in switching service providers due to contract lock-ins, switching costs across cloud platforms, and potential service disruptions.
- **Reducing Barriers to Promote Competition:**
    - Encouraging open standards and interoperability can reduce switching costs and foster competition.
    - Offering financial incentives or grants to smaller data center operators to support entry and expansion.

Key Regulatory Barriers at Various Government Levels [Question 2C]

- **Federal Level:** Complex permitting processes and stringent environmental regulations can delay projects and increase costs. OpenAI supports the goals of federal policies such as the National Environmental Policy Act while also being sensitive to the cumulative impact of multiple federal regulations on the ability of Americans to build ambitious infrastructure projects at a speed and scale that is globally competitive.
- **State and Local Level:** Zoning laws, construction permits, and local environmental impact assessments can vary widely, creating uncertainty and delays.

Existing Programs and Incentives Driving Data Center Modernization [Question 2D]

- **Government Initiatives:** Various federal programs, such as those under the Department of Energy (DOE), offer grants and loans to support renewable energy integration and energy-efficient technologies.
- **State-Level Incentives:** Some states provide incentives, rebates, or grants for data center construction, particularly in economically disadvantaged or rural areas.
- **Private Initiatives:** Partnerships with private-sector stakeholders, such as utilities or equipment manufacturers, can offer cost-sharing opportunities for infrastructure upgrades and renewable energy investments.

Actions by the Private Sector, Civil Society, or U.S. Government to Foster Data Center Growth [Question 2E]

- **Private Sector:** Invest in global partnerships, joint ventures, and strategic alliances to share resources and risks. Advocate for regulatory reforms that streamline data center construction and operations and build sustainable partnerships based on cost, supply, and regulatory efficiency.
- **Civil Society:** Raise awareness about the benefits and impacts of data center expansion, including environmental and community considerations. Engage in public consultations to support responsible development.
- **U.S. Government:** Implement policies that incentivize data center modernization, such as grants, incentives, and regulatory reforms. Promote research and development (R&D) for advanced cooling technologies and energy efficiency measures.
- **Non-Traditional Public/Private Partnerships:** Formulate longer-term plans and purchase agreements for key resources such as energy, water, labor, and talent. This requires partnerships that extend beyond traditional utilities and regulatory authorities to include arrangements at the local, state, federal levels.

Foreign Forces Pulling Data Center Market Opportunities Away from the U.S. [Question 2F]

- **Competitive International Incentives:** Foreign countries offer financial incentives, low energy costs, and access to renewable energy. Some countries also can leverage more rapid access to labor and capital that attracts data center investments away from the United States.
- **Geopolitical Risks:** Export controls, trade policies, and sanctions can affect the attractiveness of U.S.-based data centers for international customers and should be developed in a multilateral framework that doesn't hinder the competitiveness of American companies in an attempt to deny AI technology that is similar to that available from other foreign suppliers and providers.

Potential Societal Impacts of Data Center Modernization [Question 3]

- **Positive Impacts:**
    - Investing in AI infrastructure will result in tens of thousands of new jobs, significant GDP growth, a modernized grid powered by renewables like wind and solar and the construction of a state-of-the art network of semiconductor manufacturing facilities that will propel a new era of economic growth. Capital spending on AI already rivals the mainframe era of the late 1960s and the fiber optic deployment of the late 1990s – and there's an estimated $175 billion in global infrastructure funds waiting to be invested, according to an outside estimate calculated for OpenAI earlier this year. The question isn't whether that money will be spent, but where. If we don't channel it into infrastructure projects

that will support democratic AI ecosystems around the world, the funds will flow to our global competitors.
    - OpenAI recently engaged outside experts to work with us to analyze the potential job gains that would come from building 5 GW data centers in various locations across the U.S. We found that constructing and operating a single 5 GW data center could create or support about 40,000 jobs. In Texas, our analysis found that constructing a 5 GW data center could create and support more than 48,000 jobs; West Virginia could see 32,000 jobs; and Wisconsin could see 38,000 jobs.
    - Improved digital infrastructure, supporting economic growth and innovation.
    - New academic and industrial relationships that promote more efficient supply chains in the United States that integrate new fields of AI education, energy innovation, software development and hardware component manufacturing.
- **Negative Impacts:**
    - Increased energy consumption and water utilization will require adaptation and modernization. However this impact can be mitigated by clean, renewable energy, effective resource management, increased investment, and offset programs underway as discussed above.
    - Community transitions due to new land use can be offset by additional economic benefits and increased access to technology for local communities.

Supply Chain Risks, Vulnerabilities, and Threats [Question 4]

- **Critical Dependencies:** The availability of components like semiconductors, transformers, and cooling systems is crucial for data center continuity. Any shortages can disrupt operations and growth plans. Sole sources in supply chains therefore present resiliency risks.
- **Risks:** Protecting against counterfeit or untrusted components is vital to ensure data center security. Vetting suppliers, implementing rigorous auditing processes, and maintaining a robust supply chain risk management framework are essential as well as cyber and physical security procedures to mitigate risks associated with IP and Data threat or diversion.

Workforce Challenges and Opportunities [Question 6]

- **Skilled Labor Force Opportunities:** There is an opportunity for job creation and new skill development because of current labor force gaps. In particular, there is a shortage of specialized workers such as network engineers, cybersecurity professionals, electricians and construction labor. To meet this new demand, data centers will prompt new recruitment and promote retraining and retaining qualified talent. McKinsey estimates a potential shortage of up to 400,000 trade workers in the United States based

on projected data center build-out and comparable assets requiring similar skills, such as semiconductor fabrication and battery gigafactories.
- **Partnerships:** Collaborating simultaneously with a network of community colleges, apprenticeship programs, universities, research centers, and government entities to stand up a massive, localized training initiative can help develop the needed skilled workforce.

Challenges in Accessing Power for Data Centers [Question 7]

- **Reliability of Utility Power:** Data centers require reliable power, which is not always available due to grid limitations. Operators often install backup generators and explore alternative power sources, like on-site solar or natural gas generation.
- **Novel Solutions:** Data centers are experimenting with energy storage solutions, demand response programs, and microgrids to reduce dependency on grid power.

Data Center Security [Question 8]

- **Isolation in Multi-Tenant Environments:** Data center operators implement network segmentation and tenant isolation to protect data and intellectual property confidentiality, integrity, and availability across infrastructure. Isolation may be logical (software-defined) or physical (separate hardware). Isolation practices can prevent classes of vulnerabilities or architectural weaknesses that can result in cross-tenant access that could lead to data exfiltration or inappropriate access.
- **Operational Security Risks:** Data center operations and support activities including the provision of new equipment, maintenance and day-to-day operations, and end-of-life decommissioning represent potential sources of risk. These activities can involve personnel including “smart hands” interacting physically or virtually with data center or tenant infrastructure including servers. Data center operators implement measures including physical security, access controls, monitoring, two-person rules, restrictions on data-bearing devices, data destruction requirements, and other controls to prevent unauthorized intentional or incidental access.
- **Risks from Supporting Systems:** Data centers require large amounts of operational technology (OT) to support the information technology (IT) services they provide. OT systems, including HVAC, power management, access control systems, and building automation, can represent points of weakness.

Security for AI Datacenters [Question 9]

- **Single-Tenancy Presents Security Opportunities:** Reducing shared hardware and software components reduces attack surface, and with it the potential for vulnerabilities or misconfiguration to result in cross-tenant data leakage.

- **Confidential Computing Technology is Promising, but Nascent:** Data can be encrypted at rest and in transit, however opportunities to encrypt it while in use remain nascent. This is especially true for model weights: GPU confidential computing technology is promising but remains nascent, with current generation implementations having several limitations that limit their suitability for training and inference. Innovation in trusted computing and confidential computing should be encouraged. Wherever possible, AI data centers should be configured to benefit from future innovations in trusted computing, including confidential computing for GPUs.
- **Compliance with Evolving Security Standards:** Data centers should actively pursue compliance with new security standards as frameworks evolve to address unique risks. Security and compliance can be verified using third-party audits and attestation.

Respectfully,

Benjamin Schwartz

Benjamin Schwartz
Infrastructure Partnerships and Policy
OpenAI