ANTHROPIC

Build AI in America

July 2025

Executive Summary

Building AI in the United States is a national security and economic imperative. As AI systems grow more capable, the energy and computational requirements to train and deploy frontier AI are surging. Recent estimates by outside experts and our own research at Anthropic suggest that **the U.S. AI sector is on track to require at least 50 gigawatts of electric capacity by 2028,** much of which will be needed to train the world's most capable models.

America has the economic strength and technical expertise to meet these needs. But doing so will require addressing regulatory challenges—as well as supply chain, financial, and labor bottlenecks—to unlock America's innovative potential. While this task will not be easy, the United States has ample options to enable the buildout of AI infrastructure in time to sustain our AI leadership in the coming years. And the executive branch—as this report explains—can unlock this potential without new authorities for Congress, if it chooses.

Pillar 1 of our report outlines a three-part strategy for enabling frontier AI training in the United States, which will require delivering large amounts of energy to specific data centers at several locations. First, we suggest steps to accelerate permits for data centers and energy infrastructure—including by reviewing projects expeditiously and giving developers the option to build on federal lands, alongside private lands, when doing so aids permitting. Second, we propose actions to speed up targeted transmission buildouts to connect AI training facilities to power. Third, we outline steps to collaborate with utilities on accelerating interconnections to the electric grid, with options for requiring certain interconnections if necessary for national security.

As for spurring AI's wider deployment across the economy, many options in Pillar 1 will be useful, too. But diffusing AI nationwide will require a broader-based effort to unlock energy and data center buildouts around the country. To this end, Pillar 2 of our report outlines further steps to ease permitting and power procurement barriers, invest in energy infrastructure, and strengthen critical workforces and supply chains in all regions of the country.

Key facts about the AI infrastructure buildout:

- The United States is on track to require at least **50 gigawatts** of electric capacity for AI by 2028 to sustain its AI leadership.
- AI's power needs are split across two kinds of workloads. **AI training** is the process of developing AI models. We expect training models at the frontier to require data centers with **5-gigawatt capacity** in 2028. **AI inference** refers to using AI models and requires powering a broad-based network of data centers around the country.
- Bringing online the infrastructure to support these workloads requires building **data centers, generation resources, and transmission infrastructure.**
- In the United States, **building this infrastructure faces yearslong delays** from overlapping federal, state, and local permits; regulatory approvals for transmission lines; and delays interconnecting projects to the grid.
- China is rapidly building AI infrastructure and brought **over 400 gigawatts** of power onto its grid last year, compared to **several dozen gigawatts** in the United States.

Pillar 1: Building Large-Scale AI Training Infrastructure:

- **Lease DOD/DOE lands as options for data center construction** nearby BLM lands available for power procurement, avoiding the need for state and local zoning.
- **Accelerate reviews under the National Environmental Policy Act (NEPA)** through, among other things, a "programmatic" review of data centers' environmental impacts.
- **Leverage DOE transmission-partnership authorities** to avoid lengthy state reviews.
- **Support utilities on interconnection reforms,** with options available to require interconnections if needed for projects critical for national security.

Pillar 2: Building Broad-Based Infrastructure for AI Innovation Nationwide:

- **Accelerate geothermal, natural gas, and nuclear permitting,** including by reducing overlapping NEPA reviews that apply to single projects and developing "categorical exclusions" from NEPA.
- **Establish National Interest Electric Transmission Corridors** to speed permitting for transmission infrastructure in areas of key data center growth.
- **Strengthen domestic production of critical grid components and turbines** through loan and loan guarantee programs and building strategic reserves of these products.
- **Support training and entrepreneurship programs** for critical energy workers.

Introduction

America is the world's AI leader. But our leadership faces a test. As AI systems grow more capable, with profound implications for national security, our economy, and society, do we have what it takes to build AI in America? Or will America offshore the requisite infrastructure—and jobs that come with it—letting others lead in building our world's AI future?

Some argue that the United States, try as it might, simply cannot bring online the energy and computing infrastructure needed for AI leadership several years from now. Regulatory constraints and construction roadblocks, the thinking goes, make it impossible. At Anthropic, we disagree. The United States has ample options to speed up an infrastructure buildout that lets us train the world's largest AI models and use AI across our economy. This report details how the executive branch could unlock this potential without new authorities from Congress.

The core challenge is that building and using advanced AI systems requires increasing amounts of computing power and electricity. Publicly available data—as well as Anthropic's internal research—indicate that AI models' capabilities continue to grow predictably as the compute and energy used to train them scales in size. All told, we anticipate the U.S. AI sector is on track to require at least 50 gigawatts (GW) of electric capacity by 2028—a substantial portion of which will be dedicated to training several of the world's most capable models.

Experts disagree about the exact estimates. But if these numbers are approximately right, then meeting AI's power needs over the next few years will be a significant challenge. The root cause is that America's energy ecosystem isn't equipped for such rapid growth—and for years hasn't been accustomed to it. For the last two decades, U.S. electricity demand has grown an average of less than 1% per year, but AI is estimated now to be causing growth several times faster.[1] America has the economic strength and technical know-how to meet this demand. But bringing online the infrastructure needed for AI innovation—data centers, along with new power generation and transmission facilities—is facing yearslong delays from lengthy, patchwork regulatory processes. And further investment is needed to mitigate energy supply chain risks, including from reliance on imported foreign components; to address labor bottlenecks; and to sustain cutting-edge energy research and development.

---

[1] John D. Wilson *et al.*, *Strategic Industries Surging: Driving U.S. Power Demand*, Grid Strategies at 5 (Dec. 2024), https://gridstrategiesllc.com/wp-content/uploads/National-Load-Growth-Report-2024.pdf.

We at Anthropic believe that U.S. AI leadership is vital for national security, economic prosperity, and innovation. It is therefore important, given these challenges, to consider all available options to bring online U.S. infrastructure for training the world's most capable AI models, and for deploying AI widely across our economy.

These two activities—training and deploying AI—have different energy needs. Training models at the AI frontier will require delivering large amounts of energy to several locations. Pillar I of this report outlines a three-part strategy to meet this goal. First, we suggest steps to accelerate permits for data centers and energy infrastructure—including by reviewing projects expeditiously and giving developers the option to build on federal lands, alongside private lands, when doing so aids permitting. Second, we propose actions to speed up targeted transmission buildouts to connect AI training facilities to power. Third, we outline steps to collaborate with utilities on accelerating interconnections to the electric grid, with options for requiring certain interconnections if necessary for national security.

Many policies in Pillar 1 can help spur AI's wider deployment, too. Diffusing AI nationwide, however, will require a broader-based effort to unlock energy and data center buildouts around the country. To this end, Pillar 2 of our report outlines further steps to ease permitting and power procurement barriers, invest in energy infrastructure, and strengthen critical workforces and supply chains in all regions of the country.

To be sure, realizing this vision for America's AI buildout will be no simple feat. And exercising some options discussed here would involve significant complexity, political or economic costs, or levels of effort. Regardless, given AI's importance to national security and U.S. leadership, we believe it is critical to understand all available options, even those that are not easy to take, to inform deliberation and navigate the hardest tradeoffs.

These challenges are not without solutions. The federal government can take decisive action today that will enable the future's most powerful AI systems to be built here in America. Policymakers have the opportunity to do so through our proposals while keeping prices low for ratepayers both through regulatory authority as well as increasing generation supply. The key to doing so lies both in slashing red tape that holds back innovation at the federal level, and in partnering closely with infrastructure developers to jointly tackle other challenges. The rest of this report lays out this roadmap for building AI in America.

The Energy Demands of AI Leadership

How much energy does America need to lead in AI? Forecasting precisely is difficult, as public estimates for AI's power demand vary widely and are rapidly evolving. In 2023, one leading forecast predicted that AI data centers would, by 2028, need 14-19GW of electricity capacity globally.[2] More recent estimates give larger numbers. Consider the below examples from this year:

- In early 2024, Semianalysis—a leading AI and semiconductor research organization—forecasted the global power demand of "critical IT" to power AI at roughly **80GW by 2028** (with 56GW of this capacity in the United States).[3] This estimate excludes power used for cooling or other data center facility systems outside core compute, server, and networking functions.

- More recently, the RAND Corporation—extrapolating from trends in AI chip production—has projected that AI could require **117GW globally by 2028.**[4]

- Looking at the United States, the Department of Energy's (DOE) Lawrence Berkeley National Laboratory estimates that data centers writ large will operate using **74-132GW in 2028,** the bulk of which will be used for AI workloads.[5]

Today, the United States accounts for almost half of global data center power use.[6] Together, these numbers suggest that for the United States to keep pace with global AI growth—which it must, if it is to remain at the AI frontier—it must be prepared to operate **at least 50GW of power capacity for AI workloads by 2028.**

---

[2] Victor Avelar *et al.*, *The AI Disruption: Challenges and Guidance for Data Center Design*, Schneider Electric at 2 (2023), https://download.schneider-electric.com/files?p_Doc_Ref=SPD_WP110_EN&p_enDocType=EDMS.

[3] Dylan Patel *et al.*, *AI Datacenter Energy Dilemma—Race for AI Datacenter Space*, Semianalysis (Mar. 13, 2024), https://semianalysis.com/2024/03/13/ai-datacenter-energy-dilemma-race/.

[4] Konstantin Pilz *et al.*, *AI's Power Requirements Under Exponential Growth: Extrapolating AI Data Center Power Demand and Assessing Its Potential Impact on U.S. Competitiveness*, RAND Corporation at 2 (2025).

[5] Arman Shehabi *et al.*, *2024 United States Data Center Energy Usage Report*, Lawrence Berkeley National Laboratory at 6, 52-53 (Dec. 2024), https://eta-publications.lbl.gov/sites/default/files/2024-12/lbnl-2024-united-states-data-center-energy-usage-report_1.pdf.

[6] *Energy and AI*, International Energy Agency at 14 (2025), https://iea.blob.core.windows.net/assets/601eaec9-ba91-4623-819b-4ded331ec9e8/EnergyandAI.pdf.

Aggregate estimates like these, however, can obscure how AI's growing energy needs in fact pose two distinct challenges. These challenges arise from two types of workloads that AI operations require: **AI "training,"** and **AI "inference."**

- **AI training** refers to the process of developing an AI model. Training the most capable AI models involves feeding those models vast volumes of data, letting them learn broad patterns and relationships. Trained AI models are the foundation that underpins all manner of AI applications used to advance science, economic growth, and national security, and today, the most capable models are primarily developed by U.S. companies. These companies compete with each other—and developers abroad—to advance the AI frontier and train larger, increasingly capable models. This competition can unlock enormous opportunities for innovation; models' growing size and complexity, however, require increasing amounts of compute and energy for training.

- **AI inference** refers to the act of using already-developed AI models (or applications built upon them) for specific tasks. Using a model for any given task is nowhere near as computationally or energy intensive as training it. But the compute and energy needs of inference compound quickly the more widely AI is deployed throughout society.

Training and inference pose distinct energy challenges because their infrastructural needs differ. Today, the largest AI models are trained primarily within single data centers that contain enormous clusters of semiconductors, physically linked together to optimize the computations required for training. These data centers designed for AI training operate at enormous power capacities. For the United States to keep developing cutting-edge AI models, it must invest in large-scale AI training infrastructure in several strategic locations—with massive volumes of electricity delivered specifically to those places.

Inference, by contrast, does not need such highly concentrated electricity delivery. The challenge instead is building infrastructure to handle lighter computational workloads located physically nearby the customers seeking to use AI. The reason is that many AI use cases function better with low "latency"—that is, the lag time introduced from sending data back and forth between AI users and the data centers serving them. Inference workloads, then, are well handled by a broader, distributed network of data centers operating at lower capacities, which will enable powerful AI tools to be deployed across the economy nationwide.

How are AI's energy needs distributed across these tasks? Leading AI developers[7] and other experts[8] expect that training a single state-of-the-art AI model will require **a 5GW data center by 2028-2030.** That trajectory aligns with our own projections at Anthropic. If the United States unlocks its energy sector's potential, we anticipate using 2GW and 5GW data centers for a single AI training run in 2027 and 2028, respectively. With a handful of U.S. companies competing at the AI frontier and in need of similar capacity for their most advanced training runs, that statistic implies **20-25GW required in total for frontier AI training by 2028,** split across several locations. And many expect[9] that in the coming years **inference will use roughly as much or more compute and energy** compared to training.

Building and operating the infrastructure needed for these workloads in America will create tens of thousands of jobs, advance our economy, heighten our energy dominance, and secure our country's access to the most capable AI systems. Yet as the next section discusses, several barriers to infrastructure development make this task difficult.

---

[7] *See* Tony Samp *et al.*, *Powering AI Data Centers: Government and Industry Leaders Scramble to Develop Energy Infrastructure to Meet Growing Demand*, DLA Piper (Oct. 9, 2024), https://www.dlapiper.com/en-us/insights/publications/ai-outlook/2024/leaders-work-to-align-energy-infrastructure-with-growing-demand-for-ai-data-centers (characterizing the expectations of "[l]eading AI companies").

[8] *See, e.g.*, Tim Fist & Arnab Datta, *How to Build the Future of AI in the United States* (Oct. 23, 2024), https://ifp.org/future-of-ai-compute/.

[9] *See, e.g.*, Shehabi et al., *supra* note 5, at 50; Fist & Datta, *supra* note 8.

Challenges in the United States

Today, America is not on track to meet the energy needs of AI training or inference by 2028—when we expect infrastructure constraints may start significantly limiting domestic AI operations. Seeing why requires knowing what, exactly, it takes to build AI infrastructure. In general, this infrastructure breaks down into three types of construction projects:

- **The data center itself,** which houses the semiconductors and supporting IT equipment required to perform AI operations.
- **Electricity generation facilities** (i.e., power plants), which produce electricity that is delivered to the data center.
- **Transmission infrastructure** (e.g., transmission lines), used to transmit electricity or energy resources.

These three components, in turn, collectively face three main types of regulatory barriers, which must be cleared before bringing AI infrastructure fully online.

- **Permits.** Six main types of preconstruction permits—(a) land use and (b) environmental permits, at the federal, state, and local levels—apply individually to **each** component of AI infrastructure above. Especially onerous are multiyear state and local zoning processes, which are among the most common causes of energy project failures.[10] Also time-consuming are federal reviews under the National Environmental Policy Act (NEPA), which likewise can take years.[11] Beyond NEPA, other major federal permitting requirements stem from the Clean Air Act (CAA), Endangered Species Act (ESA), Clean Water Act (CWA), and the National Historic Preservation Act (NHPA).

- **Transmission approvals.** Building or upgrading transmission lines requires further approvals, notably from state public utility commissions (PUCs) or comparable state authorities. State PUC approvals can take years, especially for long interstate lines,

---

[10] Robi Nelson *et al.*, *Survey of Utility-Scale Wind and Solar Developers*, Lawrence Berkeley National Laboratory at 2 (Jan. 2024), https://eta-publications.lbl.gov/sites/default/files/w3s_developer_survey_summary_-_011724.pdf.

[11] *Environmental Impact Statement Timelines* (2010-2024), Council on Environmental Quality at 3-5 (Jan. 13, 2025), https://ceq.doe.gov/docs/nepa-practice/CEQ_EIS_Timeline_Report_2025-1-13.pdf; Scott Burton, *Accelerated Transmission Line Approvals for Priority Projects*, Norton Rose Fulbright (May 29, 2024), https://www.projectfinance.law/publications/2024/may/accelerated-transmission-line-approvals-for-priority-projects/.

resulting in a more than 10-year average completion time for new projects since 2005.[12]

- **Interconnection.** Utilities' approvals to interconnect facilities to the electric grid typically take 4-6 years for generation resources.[13] (States generally require separate approvals for non-interconnected, onsite power projects—that is, "behind-the-meter" generation.) Utilities and RTOs set interconnection rules that vary district by district, subject to guardrails in Federal Energy Regulatory Commission (FERC) regulations.

Beyond these above regulatory constraints, building AI infrastructure also faces notable challenges from shortages in critical technical workers (e.g., electricians) and supply chain risks with critical grid components, including transformers and circuit breakers. Lead times for domestic producers of these components can run approximately three years. Furthermore, reliance on imported components subjects developers both to "sudden stop" risks and to cybersecurity vulnerabilities, with "backdoors" periodically identified in energy components produced in China.[14]

An AI infrastructure project cannot operate at full capacity, importantly, before obtaining each applicable approval above (though interconnecting a subset of new generation resources, or building some behind the meter, can help data centers operate earlier at partial capacity). And each approval above—save interconnection approvals—must generally be obtained **before** construction can even begin. These facts make developing gigawatt-scale AI training infrastructure, specifically, so challenging for a 2028 timeframe. Construction alone for a state-of-the-art, gigawatt-scale AI data center is no simple task and could require, we anticipate, up to two years for facilities operating by 2028. Some generation facilities—such as geothermal or natural gas plants—face post-permitting construction times at least this long. **Preconstruction approvals** for gigawatt-scale AI training facilities, as a result, must be cleared by 2026 to operate by 2028. Yet many regulatory processes above—across the federal, state, or local levels—take at least this long on average to clear, and any can cause even longer delays for particular projects.

---

[12] Josiah Neeley & Devin Hartman, *State Permitting Challenges: Electric Transmission*, R Street (July 30, 2024), https://www.rstreet.org/commentary/state-permitting-challenges-electric-transmission/.
[13] Joseph Rand *et al.*, *Queued Up: Characteristics of Power Plants Seeking Transmission Interconnection as of the End of* 2022, Lawrence Berkeley National Lab at 31-32 (Apr. 2023), https://emp.lbl.gov/sites/default/files/queued_up_2022_04-06-2023.pdf.
[14] *See, e.g.*, Sarah Mcfarlane, *Rogue Communication Devices Found in Chinese Inverters*, https://www.reuters.com/sustainability/climate-energy/ghost-machine-rogue-communication-devices-found-chinese-inverters-2025-05-14/.

Notably, China—also vying for AI leadership—does not face the same set of regulatory constraints that we do. To be clear, **the United States should not adopt China's approach to permitting and building large infrastructure.** It is worthwhile to contrast the two systems, however, to underscore the importance of building expeditiously in our country and consistent with democratic values.

China's domestic infrastructure projects require numerous permits, as in the United States, but regulators complete them far more quickly. China's typical construction permitting timelines range from 3-6 months[15]—in contrast with the United States' yearslong NEPA reviews and subfederal zoning processes. One cause is laxer overall permitting standards.[16] But the bigger reason is that China's political system gives regulators the incentive and ability to streamline priority projects that advance political goals. China's national government, for example, sets top-down directives that cap municipal permitting timelines to certain numbers of days,[17] reflecting China's need for rapid building to meet national economic targets.[18] Additionally, local government revenues rely heavily on land leases and infrastructure projects[19] (although this dependence has eased some in recent years as

---

[15] *Construction Permit Requirements in China: 2025 Checklist*, Choi & Partners (2025), https://www.chinalegalexperts.com/news/construction-permit-requirements-china (noting, however, that "complex projects in major cities may require additional time").

[16] *See, e.g.*, Jihong Wang & Paul Kossof, *A Comparative Study of U.S. and Chinese Environmental Law with a Focus on the Real Estate Industry*, 50 International Lawyer 367, 376 (2017). *But see* Jianhua Xu & Jonathan B. Wiener, *Comparing U.S. and Chinese Environmental Risk Regulation*, The Regulatory Review (Dec. 20, 2021), https://www.theregreview.org/2021/12/20/xu-wiener-comparing-us-chinese-environmental-risk-regulation/ (positing that U.S. environmental regulation, while more stringent overall than Chinese environmental regulation, is less so than the conventional wisdom may suggest).

[17] *See, e.g.*, *China to Speed Up Construction Project Approvals in Push for Growth*, Reuters (Mar. 18, 2019), https://www.reuters.com/article/business/china-to-speed-up-construction-project-approvals-in-push-for-growth-idUSKCN1QZ142/; Ulrike Glück, *China Revises Key Construction and Planning Laws*, CMS (July 23, 2020), https://cms.law/en/chn/publication/china-revises-key-construction-and-planning-laws.

[18] *See, e.g.*, X. Frank Zhang, *Study Suggests that Local Chinese Officials Manipulate GDP*, Yale University School of Management (Feb. 11, 2020), https://insights.som.yale.edu/insights/study-suggests-that-local-chinese-officials-manipulate-gdp; Michael Pettis, *The Only Five Paths China's Economy Can Follow*, Carnegie Endowment for International Peace (Apr. 27, 2022), https://carnegieendowment.org/china-financial-markets/2022/04/the-only-five-paths-chinas-economy-can-follow.

[19] *See, e.g.*, Mengkai Chen & Ting Chen, *Land Finance, Infrastructure Investment and Housing Prices in China*, 18 PLoS ONE at 5-6 (2023); Major Tian, *The Role of Land Sales in Local Government Financing in China*, Cheung Kong Graduate School of Business (Sept. 3, 2014), https://english.ckgsb.edu.cn/knowledge/article/the-role-of-land-sales-in-local-government-financing-in-china/.

China's real estate market has stalled[20]). Many officials therefore skirt environmental reviews, shirk enforcement, or lack enforcement capacity to begin with.[21]

The result is that China, recognizing AI's geopolitical and economic importance, is building its own AI infrastructure at a rapid clip. Last year, China brought over 400GW of new generation capacity online,[22] and it has invested billions of dollars in data center hubs across its western regions through its Eastern Data Western Computing Plan.[23] The United States added roughly one-tenth as much generation capacity (excluding storage) in 2024.[24]

Many of these actions, if imported to the United States, would transgress core values in our country's political and legal systems, but U.S. policymakers must stay clear-eyed about our adversaries' capacity to cut corners and build quickly. They must ensure, moreover, that regulatory delays do not thwart us in building infrastructure needed to sustain our geopolitical advantages in the coming years.

The executive branch has the tools it needs to help industry obtain approvals expeditiously, bring new power onto the grid, and accelerate AI-driven innovation and job creation. The next sections discuss how it can empower AI developers to build large-scale AI training infrastructure in several discrete locations by 2028; unlock a broader-based infrastructure buildout to power inference; and resolve labor, supply chain, and investment bottlenecks.

---

[20] Tianlei Huang, *Chinese Local Governments' Reliance on Land Revenue Drops as the Property Downturn Drags On*, Peterson Institute for International Economics (July 5, 2024), https://cms.law/en/chn/publication/china-revises-key-construction-and-planning-laws.

[21] *See, e.g.*, Michael Meidan *et al.*, *China's Policy Pendulum Shifts Back Toward Environmental Protection, but Will Bureaucracy Get in the Way*?, The Oxford Institute for Energy Studies at 2 (Feb. 2024), https://www.oxfordenergy.org/wpcms/wp-content/uploads/2024/02/Chinas-policy-pendulum-shifts-back-toward-environmental-protection.pdf; Neveah Stover, *Analysis of the Environmental Regulations in China*, *in* Health, Work and the Environment, University of Iowa Pressbooks https://pressbooks.uiowa.edu/hwe-cph-3400-0a06/chapter/china/; *Land Use Rights in China: Rights and Disputes*, Choi & Partners (2023), https://www.chinalegalexperts.com/news/china-land-use-rights-disputes.

[22] Caroline Wang, *China Hit New Record of Solar and Wind Power Capacity Additions in 2024*, Climate Energy Finance at 1, 3 (Feb. 18, 2025), https://climateenergyfinance.org/wp-content/uploads/2025/02/MONTHLY-CHINA-ENERGY-UPDATE-Feb-2025.pdf.

[23] Georgia Butler, *China Has Spent $6.1bn Building Data Centers in the Past Two Years*, Data Center Dynamics (Sept. 3, 2024), https://www.datacenterdynamics.com/en/news/china-has-spent-61bn-building-data-centers-in-the-past-two-years/.

[24] *Industry Data*, Edison Electric Institute (2025), https://www.eei.org/en/resources-and-media/industry-data.

Pillar 1: Building Large-Scale AI Training Infrastructure

This section lays out a series of options to address each of the above regulatory obstacles to bringing gigawatt-scale AI training infrastructure online by 2028. Clearing the hurdles to building this infrastructure will be challenging—and some options discussed would be difficult or costly to exercise. Succeeding, moreover, will require enabling developers to handle each regulatory challenge above, as delays from any one can derail a project's timeline. But as this section shows, the United States has the tools it needs, under existing legal authorities, to enable this infrastructure's buildout in strategic locations and keep America at the AI frontier.

Policy Option Highlights

- **Lease DOD/DOE lands for data center construction nearby BLM lands available for power procurement**—effectively "swapping out" state/local zoning permits and related obstacles for NEPA review, which the federal government can accelerate.
- **Accelerate permits (including NEPA)** for data centers and energy by, among other things, pre-clearing NEPA requirements for data centers through "programmatic" review of their environmental impacts.
- Leverage DOE transmission-partnership authorities that **avoid lengthy state regulatory processes** to connect data centers to power.
- Support utilities in making interconnection reforms, with **legal authorities available as a last resort to require interconnections** critical for national security.

Leveraging federal lands alongside private lands

Finding sites suitable for gigawatt-scale AI infrastructure is a significant challenge, and developers should explore every opportunity they can. The federal government can greatly expand AI developers' optionality by allowing them, if they choose, to build large-scale AI training infrastructure—or core parts of it—on federally owned land.

Federal sites, as discussed in depth below, offer distinctive benefits that can help clear permitting, transmission, and interconnection roadblocks in a timely manner (in conjunction with other steps) and therefore raise economic efficiency. Specifically, the federal government could consider making Department of Defense (DOD) or DOE lands available for leases to build data centers in the western United States, which possesses ample land leasable by the Bureau of Land Management (BLM) for power procurement. Federal lands in western states are particularly promising given that BLM has completed extensive analyses of the impacts of using millions of acres of its western land for building solar projects, which can facilitate permitting. Western states also include most land in the country suitable for procuring hydrothermal power, which, contingent on steps to accelerate permitting, could be an important source of clean firm power for data centers in the next few years.

Leveraging these opportunities, of course, will require **careful site selection and preparation** in partnership with industry. As discussed further below, it is vital that siting choices take into account a region's transmission infrastructure, site-specific permitting characteristics (including high air quality; the absence of endangered species, wetlands, or cultural resources; and past development work), and the ability to avoid adverse impacts to surrounding communities.

In exploring options with federal lands, the United States should maintain an **"all of the above" mindset** to power procurement—one that lets developers procure power sources that best meet their needs under free market conditions. For AI training infrastructure over the next few years, sources like solar, batteries, and geothermal may prove the most economically efficient choices before advanced nuclear power comes online. Limiting developers' opportunities to procure some power sources but not others, however, risks making the U.S. AI sector less competitive in a period of global competition.

Permitting

As mentioned, the federal government can help developers permit projects efficiently by giving them the option to build AI infrastructure on federal lands with favorable siting characteristics. Federal lands offer several key benefits. Most importantly, building on federal sites avoids the need to clear state and local land use permits—an obstacle that commonly foils energy projects. It also insulates projects from other state law restrictions, while creating opportunities to build in places that minimally impact communities (and therefore also cause fewer project delays). Careful site selection, of course, is required to realize these benefits.

The tradeoff for bypassing these obstacles is that projects on federal lands must undergo environmental reviews per the National Environmental Policy Act (NEPA). But this tradeoff is highly advantageous for builders. President Trump has launched important work to simplify NEPA reviews with his Executive Order on Unleashing American Energy,[25] and a recent Supreme Court ruling[26] has helped restore agencies' discretion in managing review scope. And the Federal Government, ultimately, has ample authority to fulfill NEPA requirements expeditiously, whereas its levers to hasten state and local permitting are limited. Doing so will **create opportunities to bring new generation resources, data centers, and transmission infrastructure alike** online more quickly.

To accelerate NEPA reviews and leverage federal lands for large-scale AI training infrastructure, the executive branch can take the following steps—none of which are mutually exclusive:

- **Programmatic review (i.e., an advance review prior to site selection) of AI data centers.** Because frontier AI data centers may take 18-24 months to build, clearing their preconstruction permitting requirements by 2026 is vital for U.S. AI leadership. The government can start fulfilling NEPA requirements for data centers **today**—before picking any specific sites for development—through a "programmatic review"[27] that considers data centers' general environmental impacts. Agencies can later use this advance review to partially or fully satisfy NEPA obligations for any specific site picked for development.

- **Agency NEPA rules and national security exemptions.** Agencies can finish the work started in the Executive Order on Unleashing American Energy by publishing new final rules that let them follow more efficient procedures when implementing NEPA. Agencies' new rules should include national security exemptions—with a broader scope than NEPA policies' existing "emergency circumstances" exception—that allow for faster reviews for projects like large-scale AI training infrastructure with deep national security significance.

- **Categorical exclusions.** Where feasible, agencies should develop "categorical exclusions" that waive further NEPA review for components of AI infrastructure subject to the exclusion. Agencies should rapidly adopt other agencies' categorical exclusions following the procedures in the Fiscal Responsibility Act of 2023 as needed to accelerate infrastructure development. One especially helpful categorical exclusion may be the Department of the Air Force's exclusion for projects "similar to other actions [found] to have an insignificant impact in a similar setting" as

[25] 90 Fed. Reg. 8,353, 8,354-58.
[26] *Seven County Infrastructure Coalition v. Eagle County*, No. 23–975, slip op. at 2-3 (U.S. May 29, 2025).
[27] *See* 45 C.F.R. § 900.207.

established by past NEPA reviews.[28] BLM could explore adopting this categorical exclusion, for instance, to approve geothermal projects similar in scope to others for which it has issued a Finding of No Significant Impact.

- **"Tiering off" past environmental documents.** Agencies have discretion—especially in light of recent Supreme Court precedent—to save time by incorporating the findings of past environmental reviews into their NEPA analyses. For large-scale AI infrastructure, "tiering off" analyses such as those completed in 2024 for the data center project at Naval Air Station Lemoore[29] may help speed reviews. For solar power, the analysis in BLM's 2024 Solar Programmatic Environmental Impact Statement can aid permitting solar projects across more than 30 million acres of BLM lands in western states.[30] Similarly, BLM's 2008 Programmatic Environmental Impact Statement for Geothermal Leasing in the Western United States can help satisfy NEPA requirements for geothermal power on many BLM lands.

- **Prioritizing agency permitting staff.** All agencies have a limited number of staff qualified to complete permits, and allocating them toward large-scale AI training infrastructure projects is critical for meeting national security and economic needs.

- **Eliminating overlapping reviews for BLM lands.** Some energy projects—such as geothermal power—require separate rounds of NEPA review for BLM's decision to lease its land for power procurement and to construct the generation resources themselves. In these cases, BLM should develop a process to satisfy a single project's NEPA obligations with just one analysis.

- **Working with Congress on a narrow legislative exemption to NEPA** for AI data centers on federal sites, modeled after the bipartisan Building Chips in America Act.

---

[28] *Department of Defense National Environmental Policy Act Implementing Procedures: Appendix A: Department of Defense Categorical Exclusions* (CATEX), Department of Defense at 13 (2025), https://www.denix.osd.mil/nepa/denix-files/sites/55/2025/06/DOD-NEPA-Procedures-APPENDIX-A_FINAL.pdf.

[29] *Final Supplemental Environmental Assessment for the Proposed Construction, Operation, and Decommissioning of a Solar Photovoltaic System, Other Resilient Energy Systems, and a Data Center at Naval Air Station Lemoore, California*, U.S. Department of the Navy (Aug. 2024), https://cnrsw.cnic.navy.mil/Portals/84/NAS_Lemoore/Documents/Final-SEA-NAS-Lemoore-Energy-EUL_wFONSI.pdf.

[30] *Final Programmatic Environmental Impact Statement and Proposed Resource Management Plan Amendments for Utility-Scale Solar Energy Development*, Bureau of Land Management at ES-9–ES-10 (Aug. 2024).

While NEPA is the biggest permitting challenge with federal lands, policymakers should also consider—and take steps to mitigate, as needed—the risks of lengthy delays from other federal permitting requirements. The optimal approach for doing so varies by permit.

- **Permits to clear by rapidly implementing standard procedures.** In many cases, standard consultation and implementation processes for permits under the CAA (specifically, New Source Reviews, the CAA's main preconstruction permit) and ESA may be sufficient for permitting projects by 2026. The government should proceed to complete these permitting processes as quickly as possible on federal sites it makes available for AI infrastructure. For the CAA, this work could include advance coordination and technical assistance from the Environmental Protection Agency (EPA) to the state authorities implementing CAA requirements for any relevant data centers. If needed to meet 2028 AI training needs, however, the ESA includes a national security exemption available after the completion of statutorily required consultations. To this end, policymakers could also explore the CAA's exemption for federal contracts, loans, and grants as possibly applicable to data centers located on leased federal land.[31]

- **Permits to manage through careful site selection.** Permits under the CWA and NHPA can pose longer delays and may lack national security exemptions that clearly apply to AI infrastructure. Many sites, however, may not require lengthy permits under these laws—particularly with data centers' most advanced cooling technologies no longer relying on water cooling, which facilitates siting away from wetlands. If making federal lands available for AI infrastructure, the government should factor the applicability of these permits into site selection. (Careful siting is vital to minimize ESA delays as well.)

Transmission

America's strained and fragmented transmission infrastructure holds it back from bringing online large industrial projects, like gigawatt-scale data centers.[32] Yet in recent years, line construction has stagnated. Just 55 miles of high-voltage lines were built in 2023, compared to a 2010-2014 average of 1,700 miles annually.[33] State approvals are a major

---

[31] 42 U.S.C. § 7606(d).
[32] *See, e.g.*, Janelle Conaway & Barbara DeLollis, *How Can the U.S. Speed Up Energy Transmission Projects? Stop and Listen, Researchers Say*, Harvard Business School (Sept. 26, 2024), https://www.hbs.edu/bigs/how-can-the-us-speed-up-energy-transmission-projects; Eric Hirst, *U.S. Transmission Capacity: Present Status and Future Prospects*, U.S. Department of Energy at 49-51 (June 2004), https://www.energy.gov/oe/articles/doe-transmission-capacity-report.
[33] Nathan Shreve *et al.*, *Fewer Miles: The US Transmission Grid in the 2020s*, Grid Strategies at 4 (July 2024),

reason for project delays, as state regulators typically bear responsibility for approving lines and navigating politically complex siting decisions, even for interstate projects.[34]

The prospect of yearslong delays in state-level transmission approvals affects America's ability to build infrastructure for frontier AI training over the next few years. Transmission lines let data centers connect to offsite sources of power. The quality of a region's transmission resources should thus factor heavily into site selection. It is unlikely, however, that any given site will come with adequate transmission capacity to deliver gigawatts of additional power to an AI data center. We therefore believe the federal government should understand options to accelerate targeted transmission upgrades and buildouts, if needed for projects that ensure U.S. AI leadership and have national security significance.

The surest option is for the federal government to exercise federal authority for siting and approving lines, in lieu of state regulators. Taking this step can let federal policymakers ensure that mission-critical projects proceed expeditiously. With careful siting choices, the government could limit any use of these authorities for short, intrastate lines spanning several dozen miles—not hundreds of miles of interstate lines. And by approving siting at the federal level, the federal government can take steps to ensure that **the costs of transmission buildouts and upgrades are allocated to developers,** not electricity ratepayers. The executive branch has two options to take this approach in a timeframe viable for 2028:

- **Siting on federal lands.** Transmission projects that traverse only federal lands are not subject to transmission approvals at the state or local levels.[35] By making available federal lands to site data centers and generation resources, then, the government could open pathways for developers to plan routes on federal lands that require minimal state review, and for which DOE may approve cost arrangements. The western United States, for example, has significant amounts of BLM lands within 15 miles of existing or planned transmission lines[36]—creating opportunities to bring new power to the grid with only targeted transmission builds. Transmission projects on federal lands would require NEPA review, but the government can

---

https://cleanenergygrid.org/wp-content/uploads/2024/07/GS_ACEG-Fewer-New-Miles-Report-July-2024.pdf.

[34] Neeley & Hartman, *supra* note 12.

[35] Liza Reed, *Transmission Stalled: Siting Challenges for Interregional Transmission*, Niskanen Center at 7-8 (Apr. 2021), https://www.niskanencenter.org/wp-content/uploads/2021/04/Reed-Transmission-Brief-April-2021-2.pdf.

[36] *See Final Programmatic Environmental Impact Statement and Proposed Resource Management Plan Amendments for Utility-Scale Solar Energy Development*, *supra* note 30, at ES-9–ES-10.

complete these reviews efficiently using the steps described above (including by leveraging existing categorical exclusions that can apply to transmission projects).

- **DOE transmission partnership authorities.** DOE possesses several powerful statutory authorities that let it partner with private transmission developers on projects to plan, upgrade, build, operate, or finance transmission lines. These authorities—the most well-known of which is Section 1222 of the Energy Policy Act of 2005—can empower DOE to approve lines in lieu of state regulators. As past legal analysis by DOE explains,[37] Section 1222 confers these powers because the Energy Policy Act of 2005 does not waive DOE's sovereign immunity, required to give state regulators jurisdiction. While Section 1222 is only available in western states where the Western Area Power Administration or Southwestern Power Administration operate, another provision of law—Section 40106 of the Infrastructure Investment and Jobs Act—confers similar transmission partnership authorities on DOE.

To implement these steps, DOE should consider **launching a public solicitation** for proposals to build transmission lines, in partnership with DOE, that serve large-scale AI training infrastructure on federal sites. A public solicitation process—which would allow DOE to vet proposals' public impacts and consistency with regional transmission plans—may be an important step to activate some effects of Section 1222 described above. DOE should further consider leveraging its credit lines with the Department of Treasury authorized by Section 40106 and by Section 402 of the American Recovery and Reinvestment Act—a related authority—to assist with financing transmission lines built for large-scale AI training infrastructure.

Interconnection

Building large-scale AI training infrastructure by 2028 will likely require interconnecting at least some of the power resources serving AI data centers. First, the option to interconnect significantly expands available siting opportunities. Regulatory and physical constraints make it challenging to build 5GW of generation contiguously with a data center in a single location, at least by 2028, as required for a fully "behind-the-meter" project. Even America's largest nuclear plants achieve a maximum capacity of roughly 4-4.5GW,[38] and in any event, building new nuclear plants is impractical to meet AI's 2028 power needs. Additionally,

---

[37] *Summary of Findings in re Application of Clean Line Energy Partners LLC Pursuant to Section 1222 of the Energy Policy Act of* 2005, U.S. Department of Energy at 15-21 (Mar. 25, 2016), https://www.energy.gov/sites/prod/files/2016/03/f30/Summary%20of%20Findings%20Plains%20%20Eastern%20Clean%20Line%20Project%203-25-2016%20FINAL.pdf.

[38] Slade Johnson, *The United States Operates the World's Largest Nuclear Power Plant Fleet*, U.S. Energy Information Administration (Apr. 24, 2025), https://www.eia.gov/todayinenergy/detail.php?id=65104.

interconnecting AI infrastructure provides important stability and reliability benefits; simplifies construction work; and reduces financial risk, while benefitting communities, by letting developers transmit any excess power back to the grid.

Meeting AI's 2028 power needs will not, of course, require interconnecting new resources overnight—infeasible for utilities. Instead, it will mean bringing down interconnection times from 4-6 years to 2-3 years. Since utilities handle interconnection procedures, any progress in this direction will depend, in practice, on steps by utilities to accelerate operations or boost efficiency. Typically, interconnection processes involve three distinct stages: (1) queues as applications wait to be evaluated, (2) resilience and reliability testing, and (3) completion of grid upgrades required by evaluations. And **utilities can undertake vital work to accelerate each stage** of the process. Options available to utilities include the following:

- Allowing industrial projects to interconnect more quickly if they agree to limit electricity consumption on days of peak grid use when they might cause grid strain.
- Adoption of automated software tools—including AI—to conduct resilience tests themselves more quickly.
- Allowing auctions of positions in the interconnection queue. This process could involve utility studies of requisite grid upgrade costs prior to the auction, and proceeds could remunerate both generator applicants who forego queue positions as well as utilities.
- Methodology reforms that allow for batching more projects together for simultaneous review in resilience tests.
- Establishing protocols that allow select projects critical for national security, such as large-scale AI training infrastructure, to receive prioritized review in the interconnection queue.
- Allowing auctions of available grid capacity.
- Conducting advance studies of regional energy demand growth (especially associated with AI infrastructure on federal sites) and the resilience implications of interconnecting generation projects to serve that demand. These analyses would reduce the need for detailed studies in the interconnection application process.

The federal government has an important role to play in **helping utilities undertake reforms like these.** The federal government should coordinate with utilities to help them anticipate where large-scale AI training infrastructure will be located—especially for infrastructure located on federal lands—and factor utilities' input into any federal siting decisions. It can also provide technical assistance to help utilities implement changes. And it could explore legislation with Congress that provides greater support to utilities on significant operational reforms.

If needed as a last resort, however, policymakers should understand options for requiring interconnections needed for projects vital to national security. In particular, the executive branch has authorities that could create a legal requirement for utilities to interconnect such resources. Such a legal requirement—issued well in advance of its effective date—could help propel procedural changes that lower interconnection timelines to 2-3 years only for large-scale AI training infrastructure. We recognize that **exercising any option like this could be challenging** operationally, politically, or otherwise. These options should be considered, just as with other options in this report, to inform deliberations on complex national security questions.

One option to achieve this goal is **the Title I "priority" and "allocation" authorities of the Defense Production Act (DPA).** Section 101 of the DPA lets the President prioritize or allocate goods and services in the civilian market as necessary or appropriate to support the national defense, subject to certain conditions. These authorities can be applied to electric-grid interconnections—a service that utilities provide to grid customers and generation owners, as defined by an executed interconnection agreement. American leadership at the AI frontier is vital for the United States and its allies to maintain democratic countries' military advantages. We believe, therefore, that under some circumstances the President could be justified in using the DPA to require timely interconnections for frontier AI training infrastructure—only if that infrastructure required the interconnections, and only if no other means of providing them were available.

Another option to consider for this purpose is **Section 202(c) of the Federal Power Act.** This provision, when read in conjunction with the Department of Energy Organization Act, allows the Secretary of Energy to require temporary interconnections of resources needed to address "an emergency [that] exists by reason of a sudden increase in the demand for electric energy."[39] The Secretary of Energy could explore whether these conditions exist given the national security implications and energy needs of frontier AI training. This authority would require the Secretary, however, to reissue interconnection requirements every 90 days for the emergency's duration.

[39] 16 U.S.C. § 824a(c)(1); 42 U.S.C. §§ 7151(b) & 7172(a).

Pillar 2: Building Broad-Based Infrastructure for AI Innovation Nationwide

Large-scale AI training infrastructure—though necessary for U.S. AI leadership—is not itself sufficient to capture the full innovation benefits from this technology. AI's capacity to advance science, productivity, growth, and security depends on its deployment across the economy. The diffusion of AI, in turn, requires a broad-based network of data centers and supporting energy infrastructure to accommodate rising inference needs. Many options above, such as reforms to NEPA, can accelerate this broader infrastructure buildout, too. But additional pro-innovation policies will unlock incremental power builds around the country needed to capture AI's benefits.

The Trump Administration has already made progress in unlocking U.S. energy development. The Administration's announced steps to accelerate NEPA reviews will make it easier to permit projects around the country. Additionally, President Trump signed an ambitious executive order to target developing 300GW of new nuclear power by 2050, as well as 18-month timeframes for the Nuclear Regulatory Commission (NRC) to process applications for new reactors.[40] We believe there are still more opportunities to support building the physical infrastructure needed to maximize America's AI innovation potential. We discuss some key options below:

Permitting

- **Accelerate geothermal permitting.** Roughly 40GW of hydrothermal power—accessible via tried-and-tested geothermal technologies—is available in the United States, primarily on western BLM lands.[41] But as discussed in the previous section, geothermal projects at BLM sites face a "double NEPA" challenge, with distinct rounds of reviews required at the lease and powerplant construction stages. Beyond combining these reviews, as previously recommended, the federal government could seek to obviate them for some projects with a "programmatic review" of geothermal projects that satisfies NEPA requirements in advance. The Department of Interior (DOI) could also develop a categorical exclusion for geothermal powerplants themselves, which involve fewer environmental hazards than many other power projects.

---

[40] 90 Fed. Reg. 22,587, 22587-88.

[41] Robert W. Sweeny & Noah Gordon, *Geothermal Energy and U.S. Competitive Advantage: Drill, Baby, Drill*, Carnegie Endowment for International Peace (Mar. 13, 2025), https://carnegieendowment.org/research/2025/03/geothermal-energy-and-us-competitive-advantage-drill-baby-drill.

- **Continue improving nuclear permitting** by accelerating implementation of the Trump Administration's nuclear executive orders and the requirements of the ADVANCE Act of 2024.

- **Establish a nationwide permit for AI data centers**—as well as permits for other appropriate energy infrastructure—by the U.S. Army Corps of Engineers (USACE) to accelerate federal wetlands reviews under the CWA. USACE should explore opportunities to establish a nationwide permit prior to completing its next five-year reauthorization of existing nationwide permits. In tandem, USACE should prioritize allocating regional permitting staff to AI infrastructure projects seeking wetlands permits.

- **Develop additional categorical exclusions**—beyond those discussed in the previous section—useful to advance AI data center development, such as exclusions for fiber-optic cables. Agencies should adopt such exclusions developed by other agencies per the Fiscal Responsibility Act of 2023 where applicable.

- **Ensure uniform best practices on permitting and leasing across BLM regional offices.** Today, the speed of these processes can vary significantly by office—with some, for example, comparatively faster at executing leases but slower at evaluating permits, or vice versa. DOI should proliferate best practices in each area across offices including, as appropriate, via technical assistance, information exchanges, interoffice educational programs, staffing support, and engaging state permitting authorities. To gauge efficiency gains—particularly in western states with abundant BLM land—DOI should establish metrics tracking the time required for each office to perform these tasks.

- **Introduce flexibilities to backup generator requirements.** The EPA, for example, should consider updating its guidance[42] to allow data centers to operate backup generators more easily in nonemergency periods, when doing so would assist with interconnection or otherwise advance project development, consistent with grid reliability.

---

[42] For an overview of relevant guidance from the EPA and its applicability to data centers, see James Dolphin *et al.*, *New EPA Guidance Clarifies When Data Centers and Other Operators May Utilize Emergency Backup Generators to Support Local Power Supply*, Kirkland & Ellis (May 12, 2025), https://www.kirkland.com/publications/kirkland-alert/2025/05/new-epa-guidance-clarifies-when-data-centers-and-other-operators-may-utilize-emergency-backup.

- **Explore updates to EPA policies** providing for states to include national security exemptions to New Source Reviews—the CAA's main preconstruction permit—in their CAA implementing regulations. The EPA could also explore further steps to help ensure that states revise their implementing regulations as needed to attain air-quality standards as expeditiously as possible.[43]

Transmission and interconnection

- **Publish bold targets for miles of transmission lines built and generation resources interconnected** annually, including specific goals for emerging energy technologies. These targets would build on the nuclear-power objectives that this Administration has already set. For example, the National Energy Dominance Council might consider setting national goals of at least 2000 miles of annual transmission buildouts—slightly above the average a decade ago—and 1GW of new baseload power deployed in every western state by 2026.

- **Establish National Interest Electric Transmission Corridors (NIETCs) in areas of significant data center growth,** including any federal sites with large-scale AI training data centers. In areas designated as NIETCs, DOE wields substantial authorities to accelerate permits across federal, state, and local levels that impede transmission development. Because designating NIETCs requires DOE to undertake lengthy transmission studies—as well as performing a NEPA analysis—NIETCs are unlikely to assist with transmission builds on a 2028 timeframe. The executive branch could work with Congress, however, to shorten some of these requirements, create a categorical exemption for NIETC designations, or affirm that NIETCs provide DOE the authority to intervene in states' permitting processes when those processes do not consider projects' national security benefits accruing to other states.

- **Allow utilities and developers to propose NIETCs directly to DOE,** consistent with the Federal Power Act's directive to consider "recommendations from interested parties" in NIETC designations.[44] Doing so would make DOE's designation process more responsive to market signals, and DOE can further accelerate designation by allowing the recommending parties to conduct NEPA analyses at the time of proposal. Additionally, DOE could automatically start a NIETC study in any region where contracts emerge for "merchant" transmission projects—indicating unmet market demand for transmission capacity.

---

[43] *See* 42 U.S.C. § 7410(a)(2)(H)(i).
[44] 16 U.S.C. § 824p(a)(2).

- **Identify opportunities for fast interconnections** by using DOE information-collection authorities[45] to require utilities and large grid customers to report underutilized points of interconnection, surplus interconnection service, and related information about interconnections they operate. DOE could make this information available to developers seeking to bring new power online and collocate it near existing facilities. Additionally, DOE could use these authorities to establish requirements for utilities to report information about transmission congestion, which aids transmission planning.

- **Require utilities to deploy grid-enhancing technologies (GETs) and other grid efficiencies** via FERC rulemaking, rather than merely encouraging GETs. The executive branch could also collaborate with Congress on legislation in this area.

- **Simplify FERC regulations for or otherwise encourage "merchant" transmission projects** to allow for greater flexibility in financing transmission lines needed to serve new AI data centers.

Financing AI innovation

- **Transmission financing.** Two of DOE's transmission partnership authorities discussed above—Sections 402 and 40106—give DOE $3.25 billion and $2.5 billion, respectively, in credit lines for lending to transmission developers. DOE should increase use of these authorities for AI infrastructure, specifically targeting projects in need of additional financing to support AI training or deployment. As part of this effort, DOE should explore the legal availability of selling debt instruments it creates from its balance sheet to third-party buyers—thereby replenishing its credit line and enabling more lending, on a repeatable basis.

- Launch a **loan-guarantee program to support domestic developers of critical grid components** making capital expenditures or scaling up worker training to help meet the need for more products and materials. Loan guarantees could offload risk and expand access to financing, in particular, for producers of transformers and circuit breakers wary of growing capacity owing to concerns that AI's surging power needs might flag.

- Expand **loan and loan-guarantee offerings for nuclear technology**, or other forms of support, to mitigate risks from cost overrun—including by working with Congress on bipartisan legislation that supports nuclear financing.

---

[45] 15 U.S.C. § 772.

- Scale up DOE grants and other **support for demonstration projects** that advance the state of the art in next-generation geothermal technologies,[46] hydropower, and advanced nuclear technology. This vital support for innovation creates jobs and shortens the time to commercialization for energy technologies of the future.

Strengthening supply chains and workforce

- Establish **strategic reserves of critical grid components** by guaranteeing purchases of transformers, circuit breakers, combustion turbines, organic ranking cycle (ORC) turbines, and other equipment when prices fall below a certain level. Building reserves this way will simultaneously offer price support to producers while ensuring that the United States maintains supplies from which it can draw in a crisis.

- Coordinate with industry to **develop ground-truth supply and demand forecasts** for critical grid components to assist with industry planning and inform capital expenditures. The Department of Homeland Security (DHS), DOE, and the Department of Commerce could collaborate on collecting necessary industry data and developing forecasts.

- Impose **cybersecurity requirements on imported energy products,** such as substations, that could carry “backdoor” vulnerabilities. DOE, the North America Electric Reliability Corporation, and DHS could develop standards collaboratively with industry. One option for imposing them could be the Federal Communications Commission’s authorities that allow it to set standards and requirements for products or components that emit radio frequencies. The government could also explore using the International Emergency Economic Powers Act (IEEPA)—which authorizes the regulation of foreign transactions that pose national security risks—to block imports of critical components unless they meet certain cybersecurity standards.

- Expand **financial support and technical assistance for employers offering apprenticeships** for electricians, mechanical engineers, and construction roles critical to building AI infrastructure. Support could be routed through the registered apprenticeship program or other appropriate mechanisms.

---

[46] *Enhanced Geothermal Systems (EGS) Pilot Demonstrations*, U.S. Department of Energy, https://www.energy.gov/eere/geothermal/enhanced-geothermal-systems-egs-pilot-demonstrations.

- Support **entrepreneurship programs for energy workers**—from oil and gas workers to solar power engineers—in partnership with industry, nonprofit organizations, and higher-education institutions that have leading geosystems engineering or other energy programs. Programs that provide energy entrepreneurs with capital, technical assistance, mentorship, and market access will pave the way to future energy breakthroughs that secure America's energy dominance.

- **Engage higher-education institutions,** including community colleges, on other opportunities to expand training and educational offerings in these areas. These institutions could do so in partnership with companies involved in construction of large-scale AI infrastructure.

Conclusion

Building the infrastructure required for AI leadership is no simple task. Beyond technical design and construction challenges, bringing AI infrastructure online requires navigating a wide range of regulatory, supply chain, financial, and labor obstacles—any one of which can delay a given project by years. Developers must be able to address **each one** of these constraints to bring AI infrastructure projects online fully. Addressing them, moreover, takes considerable time and planning, and the strategies for doing so interact with complex siting choices that occur early in projects' lifecycles. The United States must therefore act **today** to unlock our country's potential to build, achieve full-spectrum energy dominance, and secure our AI leadership for the future.

The executive branch has options to do so if it chooses. Some options, of course, may be challenging to implement for operational, political, or other reasons—or only appropriate as a last resort. We at Anthropic believe, however, that it is vital for federal policymakers to understand the full range of options available as they navigate hard tradeoffs. The actions that America takes today will have profound consequences for national security and technological, economic, and energy leadership for years to come.

Appendix

To advance deliberation further on policy options for large-scale AI training infrastructure, below are notional timelines that could serve as targets for key milestones in data center, generation, and transmission projects—contingent on implementing policy options in this report. The targets below are **not** forecasts. The time needed to permit and construct any given project, importantly, will be highly site specific. Particular processes may require more time, for some projects, than the rough targets suggest below; others may be able to proceed more quickly. Developing any project expeditiously also hinges on effective site selection. Still, we believe that even notional targets like these can assist policymakers and developers in evaluating different options for building AI infrastructure—and in identifying processes involved with certain projects that may proceed more quickly, or that may require more time.

Selected factors to consider in site selection:

- Existence of past environmental documents or development on the site that can lead to shorter permitting timelines.
- Minimal adverse impacts from development on communities.
- Adequacy of regional transmission infrastructure and scope of required transmission development.
- High air-quality standards for purposes of the Clean Air Act.
- Absence of major permitting challenges related to wetlands (for the Clean Water Act), endangered species (for the Endangered Species Act), and historical or cultural artifacts (for the National Historic Preservation Act).
- Topographical and other technical characteristics of land conducive to the project being developed.

For purposes of illustration, the notional targets below focus on milestones for completing construction and major permits and approvals needed to build projects on federal lands, based on publicly available information about different projects. They assume that policy options in this report are implemented successfully, that site selection has incorporated the factors above, and that supply chain or financing challenges do not significantly delay projects.

Notional targets for AI training data centers:

- **July 2025 – Dec. 2025:** Select DOE and DOD sites suitable for data centers, solicit project proposals via competitive leasing process, and select applicants.

- **July 2025 – June 2026:** Programmatic review of data centers' environmental impacts to help satisfy NEPA requirements.
- **By June 2026:** Any additional needed site-specific analyses or processes to satisfy NEPA requirements.
- **By early 2028:** Construction of data centers for frontier AI training.

Notional targets for solar projects:

- **July 2025 – December 2025:** Competitive leasing processes for solar projects on BLM lands able to deliver power to sites eligible for AI training data centers.
- **By June 2026:** Complete NEPA reviews for construction of solar projects
    - *Illustrative options to advance this goal:* "Tiering off" past environmental reviews, including BLM's programmatic solar EIS; applying or developing appropriate categorical exclusions, such as DOE's exclusion for solar projects; prioritizing federal permitting staff.
- **July 2026 – Jan. 2028:** Construction of solar plants and storage facilities.[47]

Notional targets for natural gas projects:

Note: As mentioned above, the notional targets below do not factor in lead times for turbines or other components.

- **July 2025 – Dec. 2025:** Competitive leasing processes for natural gas projects on BLM lands able to deliver power to sites eligible for AI training data centers.[48]
- **Jan. 2026 – Late 2026:** Complete NEPA reviews and Clean Air Act permits required for construction of natural gas plants.
    - *Illustrative options to advance this goal:* "Tiering off" past environmental reviews; prioritizing federal permitting staff; updates to agencies' final rules

[47] Depending on complexity, solar projects can take as little as 8 months to construct or an upper bound of 18-24 months. For a range of estimates, see, for example, *Understanding the Solar Project Development Process Steps*, Urban Grid (Jan. 16, 2019), https://www.urbangridsolar.com/understanding-the-solar-project-development-process-steps/; Mark Richardson, *How Long Does It Take to Build a Solar Farm?*, U.S. Light Energy (Nov. 10, 2023), https://uslightenergy.com/how-long-does-it-take-to-build-a-solar-farm/. The notional targets above conservatively assume longer development timelines consistent with utility-scale projects. Construction of storage facilities—simple or complex—typically involves shorter timeframes. *See, e.g.*, *What Is Utility-Scale Energy Storage?*, Arevon (Dec. 19, 2024), https://arevonenergy.com/news/blog/what-is-utility-scale-energy-storage/.
[48] This notional target aligns with BLM's objective of executing natural gas leases within 6 months from start to finish, including scoping and environmental analysis. *See Oil and Gas Leasing—Land Use Planning and Lease Parcel Reviews*, IM 2025-028, Bureau of Land Management (May 8, 2025), https://www.blm.gov/policy/im-2025-028.

implementing NEPA; advance engagement with state air-quality permitting agencies.

- **By late 2026:** Preconstruction design work completed for natural gas plants.
- **By late 2028:** Construction of natural gas plants (with exact target construction timelines varying by plant size).[49]

Notional targets for geothermal projects:

Note: Geothermal projects have especially uncertain timelines, as key steps such as resource exploration and confirmation may occur without full visibility into a site's potential for geothermal power. Permitting and financing delays, traditionally, have also extended project timelines significantly.[50] At the same time, innovations in drilling and fracking technology can help accelerate certain drilling operations for geothermal projects. The notional targets below factor in typical timelines for these operations, but different projects will proceed at different paces.

- **July 2025 – Dec. 2025:** Competitive leasing process for geothermal projects on BLM lands able to deliver power to sites eligible for AI training data centers.
    - *Illustrative options to advance this goal:* "Tiering off" BLM's programmatic review of geothermal leases; prioritizing federal permitting staff; adopting and applying a categorical exclusion for projects similar to others that have received a Finding of No Significant Impact, or that will undergo subsequent environmental reviews.
- **Jan. 2026 – Late 2028:** Resource exploration, confirmation, production drilling, and power plant construction for geothermal projects (assuming typical operational timelines for these project stages).[51]

[49] Larger natural gas plants can take 2-3 years to construct, although smaller projects can be built in shorter timeframes. *See, e.g.*, *Panda Power Temple Projects*, Bechtel (2025), https://www.bechtel.com/projects/panda-power-temple-projects/; *Panda's New Power Plant up and Running in Texas*, Factor This Power Engineering (Nov. 25, 2014), https://www.power-eng.com/coal/plant-decommissioning/panda-s-new-power-plant-up-and-running-in-texas/; (indicating a 26-month period between groundbreaking and plant operation for a natural gas plant of approximately 750MW).
[50] *See, e.g.*, Koenraad F. Beckers & Katherine R. Young, *Technical Requirements for Geothermal Resource Confirmation*, 42 GRC Transactions at 8 (2018), https://www.energy.senate.gov/services/files/9E245C4D-2082-48E6-9AB2-934362585FC0.
[51] This notional target factors in 9 months of sitework for resource exploration, see, for example, *Geothermal Exploration Best Practices: A Guide to Resource Data Collection, Analysis, and Presentation for Geothermal Projects*, IGA Service GmbH at 9 (Mar. 2013), https://www.globalccsinstitute.com/archive/hub/publications/138043/geothermal-exploration-best-practices-guide-resource-data-collection-analysis-presentation-geothermal-projects.pdf, between 12 and 15 months for resource confirmation, see, for example, Beckers & Young, *supra* note 50, at 7-8, and approximately 12 months for drilling production wells and building the geothermal

- **By late 2027:** Any additional reviews needed to satisfy NEPA requirements for geothermal power plant construction (with previous stages of project development covered by existing categorical exclusions).

<u>Notional targets for transmission projects:</u>

- **July 2025 - Dec. 2025:** Competitive solicitation process for proposals for developers to build transmission lines needed for AI training infrastructure in partnership with DOE.
- **By June 2026:** Complete NEPA reviews and obtain siting and cost allocation approvals from DOE
    - *Illustrative options to advance this goal:* Adopting and applying existing categorical exclusions to transmission buildouts (or developing new exclusions for projects not covered); siting and building lines within existing rights-of-way subject to past development; prioritizing federal permitting staff.
- **By early 2028:** Construction of transmission projects.[52]

---

power plant itself. Production drilling timelines depend on numbers of rigs deployed and opportunities for drilling production wells in parallel. Past literature has cited 4 years total as a fast timeline for geothermal project development, inclusive of permitting and financing delays, and not incorporating recent technological developments. *See*, *e.g.*, Valerio Micale *et al.*, *The Role of Public Finance in Deploying Geothermal: Background Paper*, Climate Policy Initiative at 6 (Oct. 2014).

[52] Construction processes for transmission lines can take 12 months, particularly for smaller projects, though very large projects can require 18-24 months, depending on site conditions. *See*, *e.g.*, *SunZia Transmissions Project*, US, Power Technology (July 5, 2024), https://www.power-technology.com/projects/sunzia-transmission-project-us/ (describing a 550-mile transmission project for which approximately two years elapsed between groundbreaking and operation). This notional target assumes that site selection has ensured that only relatively short transmission projects (e.g., several dozens of miles or less) are needed for AI infrastructure projects.