CONSERVATIVE COALITION *for* CLIMATE SOLUTIONS

Considerations for Siting Data Centers on DOE Lands

An RFI response from C3 Solutions regarding the building of AI infrastructure on DOE lands.

by Ian R. Banks

May 2025

The Conservative Coalition for Climate Solutions (C3 Solutions) welcomes the opportunity to respond to DOE's Request for Information on Artificial Intelligence Infrastructure on DOE Lands—6450-01-P, published on April 7, 2025. C3 Solutions is a non-partisan 501(c)(3) think tank focused on accelerating innovation to meet America and the world's greatest energy and environmental challenges.

Vice President JD Vance recently affirmed that "there is no AI future without energy security and energy dominance" and the United States must act quickly to secure its competitive advantage in AI dominance through unbridled energy dominance. This RFI is an example of the Department of Energy setting out to do just that.

Artificial intelligence has emerged as a critical component of American national security and offers paths toward profound abundance and prosperity. It is right and good that President Trump and Secretary Wright have focused national attention and resources on building the infrastructure and culture necessary to even further accelerate American leadership in artificial intelligence. This request for information is a welcome call to utilize our great National Laboratories to support artificial intelligence infrastructure.

This response centers on the view that, for the effective advancement of American artificial intelligence leadership, the chief, if not all-surpassing consideration should be speed-to-power, that is, the time in which a data center moves from first shovel in the ground to fully powered and operational. This is an especially important consideration in this process where massive investments will be made in federally owned lands and infrastructure. By adopting speed-to-power as the core objective of the deployment phase of this initiative, projects will be incentivized to get data centers powered and operational as quickly as possible, reducing downtime and providing opportunity for tremendous gains in artificial intelligence research and training. A key way of prioritizing this metric for these projects is through milestone-based awards for achieving certain power generation goals within a specified timeline. Additionally, we affirm the suggestions from other scholars to use these parcels to establish Special Compute Zones—federally designated areas where DOE is granted preemptive siting and permitting authority—to concentrate and accelerate the deployment of co-located AI and energy infrastructure.

By adopting speed-to-power as the core objective ... projects will be incentivized to get data centers powered and operational as quickly as possible.

Below, we respond to particular questions in greater detail.

1.2 — WHAT CHARACTERISTICS OF A SITE MAKE IT MORE OR LESS FAVORABLE FOR DEVELOPMENT?

The most favorable sites for AI infrastructure development combine existing power generation footprints with completed environmental reviews, robust grid connectivity, and access to specialized technical workforce—creating "power-ready" locations that dramatically accelerate deployment timelines for advanced clean energy and computing facilities.

Energy Generation Footprints & Grid Connectivity: Sites become premium candidates when hosting or directly adjoining existing power facilities with substantial substation infrastructure, high-voltage transmission (U230kV) with verified capacity, and established rights-of-way. The critical advantage lies in rapid deployment of advanced clean firm power technologies including SMRs, microreactors, or enhanced geothermal systems. Leveraging existing switchyards, cooling infrastructure, and transmission corridors can eliminate years from development timelines, directly addressing the "speed-to-power" imperative. The value derives not from raw acreage but from pre-disturbed energy infrastructure footprints primed for next-generation power technologies.

Pre-Established Environmental Documentation: Sites with completed site-wide or programmatic Environmental Impact Statements (EISs) or recent Environmental Assessments (EAs) for comparable industrial activities provide decisive timeline advantages. This established body of regulatory-approved analysis (geological, hydrological, ecological, and cultural resource surveys) enables new AI and energy projects to tier from existing NEPA documentation or require only focused supplemental reviews. This transformation of multi-year uncertainty into predictable permitting pathways positions these sites as optimal candidates for meeting urgent compute capacity demands.

Grid Resilience with Streamlined Interconnection: While on-site firm power represents the optimal solution for AI facilities on DOE lands, access to robust existing grid infrastructure with designated high-capacity interconnection points provides essential flexibility for initial operations, redundancy, and potential clean power export. Preferred sites feature transmission networks with documented reliability metrics and where DOE can establish expedited interconnection protocols or dedicated queue processing for these nationally critical projects. This circumvents the interconnection bottlenecks currently plaguing private development and ensures power delivery aligns with AI deployment timelines.

Specialized Technical Workforce Ecosystem: The presence of a ready technical workforce with scalable capacity proves essential for both accelerated construction and sustained operations. Key personnel requirements include expertise in data center management, high-voltage electrical systems, advanced energy facility operations (particularly transferable skills for nuclear or geothermal systems), and specialized construction trades. Proximity to technical institutions, engineering-focused universities, and communities with populations possessing relevant technical backgrounds substantially enhances site viability and operational readiness.

Figure 1.

DOE Proposed AI Data Center Locations

Map of DOE-affiliated research and security sites overlaid with existing high-voltage transmission lines — highlighting co-location opportunities for firm power and AI infrastructure.

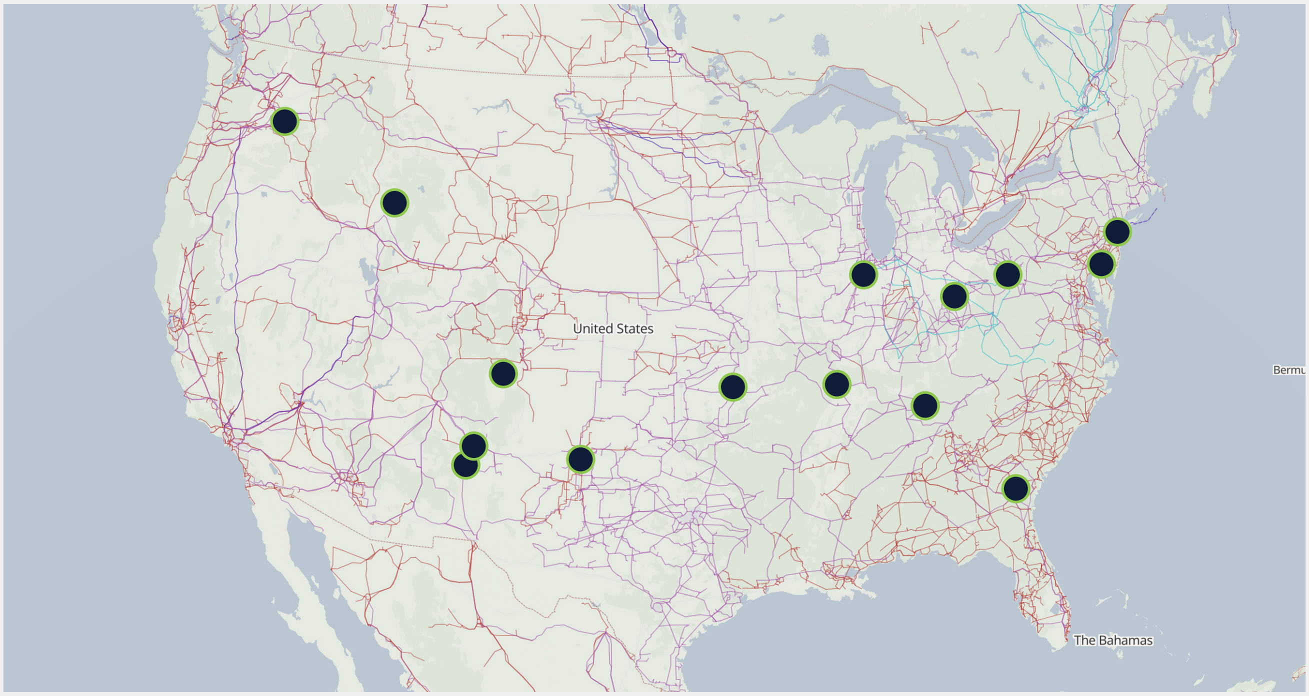

SOURCE: U.S. Department of Energy, Open Infrastructure Map

Cultures of Creative Problem Solving: The most eligible candidate sites and bids should clearly communicate a culture of creative problem solving which does not remain bound by the status quo of government contracting but seeks all available, legal, and effective means to complete the goals of the project. This creativity should include the use of Other Transaction Authorities among other novel approaches, such as seeking FAST 41 for permitting approval.

2.6 — WHAT KINDS OF ZONING (EX: SETBACKS OR HEIGHT RESTRICTIONS), LAND USE PLANNING OBJECTIVES, OR PERMITTING JURISDICTIONS ARE FAVORABLE FOR SITE CONSIDERATION?

As has been recommended by scholars from the Institute for Progress and the Foundation for American Innovation, special designations should be given to the selected parcels as "Special Compute Zones." Within these zones, a single federal entity, ideally DOE, should be granted preemptive authority for all siting, construction, and operating permits, superseding routine local zoning. Permitting decisions for both AI facilities and their dedicated energy infrastructure within SCZs must be reached within an accelerated timeframe, for example, a statutory 6-month period from complete application, with a presumption of approval if projects meet rigorous, pre-defined safety, environmental, and national security standards.

Within SCZs, using authority under the Defense Production Act DOE should swiftly establish broad Categorical Exclusions (CatEx) under NEPA for AI data centers and their associated advanced energy infrastructure (e.g., SMRs, geothermal plants, gas generators), particularly on previously disturbed or industrialized DOE lands. This recognizes an overriding national urgency and allows resources to be focused on compliance with other critical environmental statutes such as the Endangered Species Act (ESA), Clean Water Act (CWA), and Clean Air Act (CAA), which would still apply but approval should be located in a single entity, namely DOE. Furthermore, to provide essential investment certainty and prevent strategic litigation from derailing nationally critical projects post-approval, legislation should provide robust cover from dilatory judicial injunctions once all necessary permits within the SCZ framework have been issued, provided projects remain in compliance with their permit conditions. This level of certainty is crucial to safeguard the substantial capital investments required for these nationally significant AI and energy projects, preventing them from being stymied by protracted legal challenges after thorough federal review and approval.

3.1 — WHAT TYPE OF CO-LOCATED ENERGY TECHNOLOGIES ARE OF HIGHEST INTEREST IN BEING DEVELOPED WITH AI DATA CENTERS? WHAT TYPE OF SITE INFORMATION WOULD NEED TO BE PROVIDED TO INFORM USE OF A GIVEN ENERGY TECHNOLOGY (E.G., SUBSURFACE DATA, SOLAR RESOURCE POTENTIAL)?

All energy sources and technologies should be considered from bids but priority should be given to projects which utilize technologies and contractors with a proven track record of rapid speed-to-power. As discussed in the answer to 4.3, below, milestone rewards should be used to incentivize speed-to-power and ensure compliance with aggressive timelines for construction and operations. Because the cost of computing hardware is so much higher than the cost of energy generation, 24/7 power is likely necessary to make projects as cost-competitive as possible. This prioritizes firm power such as gas, nuclear, geothermal (conventional and enhanced), or renewables with storage. For renewables, given the current cost of storage and supply chain reliability concerns, priority must be given to gas, nuclear, and geothermal. Gas presents the best opportunity for rapid power deployment, notwithstanding turbine supply chain concerns. Nuclear and geothermal both offer exciting possibilities in their own right but the timelines for deployment of new generation make for longer speed-to-power timelines over gas. Where appropriate, particularly at sites where there is a demonstrated capability to rapidly deploy nuclear and geothermal energy sources, these technologies should be utilized.

4.3 — ASSUMING ADDITIONAL CAPACITY COULD BE PROCURED OR BUILT IN STAGES, WHAT ARE DESIRED TIMELINES FOR ELECTRICITY CAPACITY AVAILABILITY?

To meet the urgent power needs of AI data centers and avoid costly delays or overreliance on stop-gap fossil generation, we propose that timeline compliance should be encouraged by clearly defined milestones with rewards for adherence and greater rewards for overperformance. Milestones should be technology agnostic and should include rewards for achieving interim operational capacity such as with stop-gap fossil generation or modular power but these rewards should be less than for firm or clean power. Crucially, for projects which fail to achieve certain milestones within an appropriate time limit, any loan funding should be clawed back and in egregious cases, fees should be levied for the loss of time and the opportunity cost for another project to have been completed on the same parcel. Here is an example of an aggressive speed-to-power timeline with milestones that would be accompanied with rewards:

1. **Milestone 1 (Initial Operational Capability):** Delivery of at least 5 MW to a new data center within 6 months of project acceptance. This is in line with Super Micro and xAI's 122-day construction timeline of the Memphis supercluster, "Colossus." This could be achieved using stop-gap power generation if necessary.

2. **Milestone 2 (Firm Operating Capability):** Delivery of 50–100 MW of firm power within 12 months of project agreement, enabled by streamlined permitting on DOE lands and prioritizing sites with existing usable infrastructure or those suitable for rapid deployment of new generation sources or interconnection.

3. **Milestone 3 (Operations Expansion):** An additional 200–300 MW of firm capacity (e.g., from SMRs or EGS) online within 2 years exclusively for data centers.

4. **Milestone 4 (Full Operating Capability):** Deployment of 500–1000 MW capacity within 4 years, demonstrating a clear pathway to meet significant demand growth.

5. **Milestone 5 (Scaled Compute Zone):** Within 5-10 years, achieving 5+ GW of new power generation capacity exclusively for AI data centers.

Figure 2.

Speed-to-Power Milestones for AI Data Center Deployment

This 5-stage deployment framework illustrates the speed-to-power imperative: accelerating delivery of firm power to AI data centers on DOE lands. Each milestone reflects a critical power delivery goal—beginning with minimal operational capability within 6 months and scaling to 5+ GW over a 10-year horizon. Projects should be rewarded for early delivery and penalized for delay.

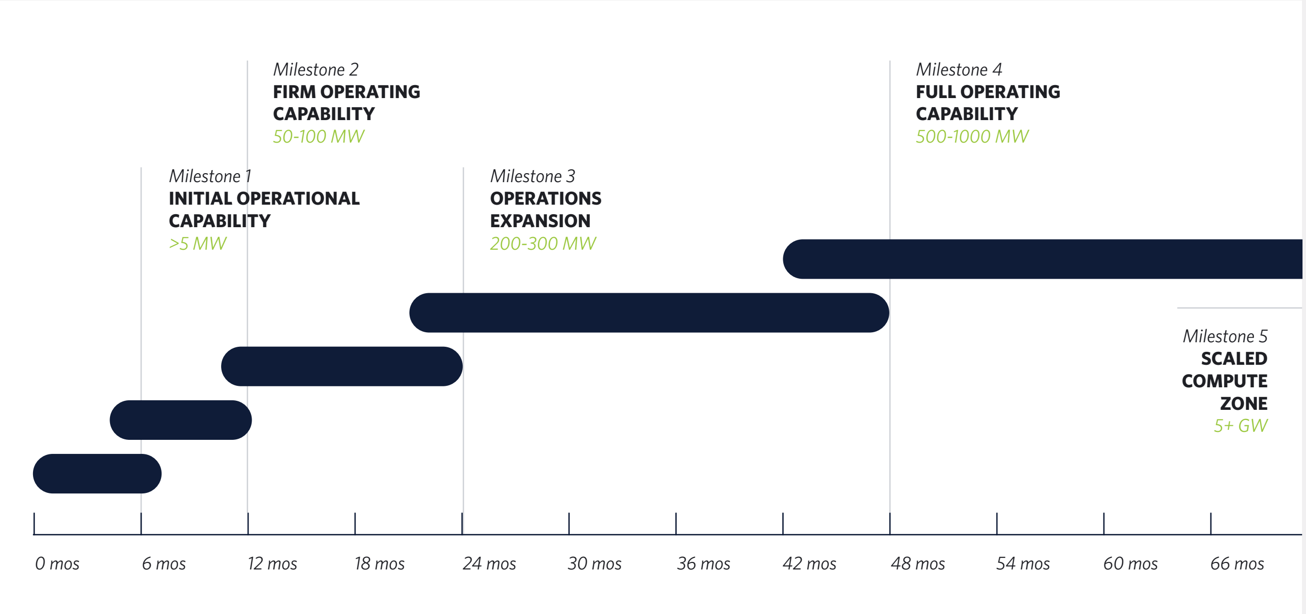

4.4 — WOULD FLEXIBLE DATA CENTER OPERATIONS BE POSSIBLE IF IT WOULD ENABLE FASTER CAPACITY INTERCONNECTION?

Maximum flexibility should be afforded to data center operators and flexible operations should not be a universal requirement for interconnection but allowing operators to choose flexible operations in return for extremely rapid speed-to-power or a greater share of renewable power generation may be in the interest of certain data center operators. The high costs of allowing extremely expensive chips to lay idle likely make flexible operations unattractive to most operators unless there is a very significant interconnection advantage. Consideration of flexible operations agreements should be allowed in bids where data center operators seek it but should not be a bargaining tool utilized by the Department instead the emphasis on speed-to-power should be overriding.

6.3 — WOULD SHARING COMPUTATIONAL RESOURCES OR PROVIDING COMPUTE CREDITS TO RESEARCHERS FROM DOE OR LOCAL UNIVERSITIES BE POSSIBLE?

Where possible and if made worth it through serious advantages from the DOE siting of the data centers, some degree of compute resource priority or credits should be given to local DOE researchers and university partners but project agreements should not be conditional on this. Specifically, a negotiated portion of AI compute (e.g., 3-5% of off-peak capacity) could be made available for National Lab and university research explicitly aimed at accelerating the commercialization and deployment of next-generation firm clean energy technologies (advanced nuclear, EGS, long-duration storage). This creates a virtuous cycle where the AI energy investments also help speed up the development of novel power solutions.

9.2 — WHAT CONCERNS EXIST WITH SUPPLY CHAIN LIMITATIONS, SUCH AS LONG LEAD TIMES ON CERTAIN POWER AND ONSITE ENERGY EQUIPMENT, AND WHAT ALTERNATIVES SHOULD BE CONSIDERED?

Severe supply chain constraints—especially multi-year lead times for high-voltage transformers, advanced switchgear, Rankine cycle turbines, and specialized nuclear components—pose a direct threat to the "speed-to-power" imperative essential for U.S. AI competitiveness. DOE should immediately employ Defense Production Act (DPA) Title I rated orders to prioritize the production and delivery of these mission-critical systems for national security-aligned AI projects. In parallel, DPA Title III investments must expand domestic capacity for long-lead equipment and upstream inputs such as grain-oriented electrical steel and precision cooling technologies.

To reduce strategic dependency, particularly on foreign-controlled sources for large transformers and thermal management systems, the U.S. must pursue secure reshoring strategies. This includes building out a Strategic Equipment Reserve, standardizing component designs for rapid scale-up, and enforcing procurement restrictions against adversarial suppliers. These actions should be matched with workforce development initiatives focused on electrical manufacturing and grid construction trades, ensuring that AI deployment timelines are grounded in realistic assessments of domestic industrial and labor capacity.

CONCLUSION

In conclusion, securing America's AI advantage hinges on making speed-to-power the overriding priority for infrastructure development on DOE lands. This means selecting power-ready sites with existing energy footprints, grid assets, and completed environmental reviews—while simultaneously enacting bold regulatory reforms through federally designated Special Compute Zones. Within these zones, DOE should be granted preemptive, time-limited permitting authority (e.g., 6-month deadlines with a presumption of approval), empowered to apply Categorical Exclusions under NEPA using Defense Production Act (DPA) authorities, and backed by statutory protections against post-approval legal delays. Execution must be anchored in aggressive, technology-neutral deployment milestones—with meaningful financial incentives for early delivery and penalties for delay—favoring firm power sources such as gas, nuclear, and geothermal. To ensure delivery is feasible, DOE must also confront energy supply chain bottlenecks head-on through DPA-rated procurement and targeted investments to expand secure domestic manufacturing of long-lead components.