

Working Paper

Could the Permian Basin Become America's Next AI Data Hub?

Gabriel Collins, J.D.¹

Baker Botts Fellow in Energy and Environmental Regulatory Affairs

This publication was produced in collaboration with Rice University's Baker Institute for Public Policy. It has not been through editorial review. Wherever feasible, this material was reviewed by outside experts before it was released. Any errors are the author's alone.

This material may be quoted or reproduced without prior permission, provided appropriate credit is given to the author and Rice University's Baker Institute for Public Policy. The views expressed herein are those of the individual author(s), and do not necessarily represent the views of Rice University's Baker Institute for Public Policy.

© 2025 Rice University's Baker Institute for Public Policy

Could the Permian Basin Become America's Next AI Data Hub?

Gabriel Collins

Computing Power is a Vital Engine of US National Power

Winning the electricity race is critical to winning the global AI race. The United States' present AI pre-eminence relies in part on a world-class tech and software ecosystem but equally, on a multi-gigawatt "digital industrial base" of massive data centers. Further rapid scaleup of this electricity-hungry physical system will determine whether the U.S. retains global AI leadership, and with it, global economic leadership – the most fundamental building block of comprehensive national power. The rapidly evolving AI competition between the U.S. and the People's Republic of China is likely to have existential consequences and the U.S. urgently needs additional computing power to retain its edge.

If the U.S. can achieve an AI-driven productivity boom it will likely cement its role as the world's largest and most technologically innovative economy. To continue leading the AI boom, the U.S. now faces an industrial and technological challenge on the scale of the Manhattan Project our forebears executed 80 years ago. This time, the investment is private but strategic facilitation will come from Washington and state capitals. U.S. government public statements thus far focus on building AI compute infrastructure in Pennsylvania and on certain federal lands, including the Idaho National Laboratory, Oak Ridge Reservation, Paducah Gaseous Diffusion Plant and the Savannah River Site.

The focus thus far omits a critical geography – the Permian Basin, which is one of the world's premier energy production areas. To that point, Texas and New Mexico both have real opportunities to attract slices of capital spending from an AI-driven infrastructure investment boom larger than anything since the Shale Revolution. For perspective, Microsoft plans to spend at least \$40 billion on AI-related data centers in the U.S. during 2025.² In contrast, Exxon, the largest U.S. oil producer, plans to spend between \$27 and \$33 billion on capital projects globally from 2025 to 2030.³

Absent decisive action, progress risks being bottlenecked and strategic advantages lost. As computing power is deployed, its innovation and productivity effects tend to compound upon each other and are multiplicative (i.e., exponential) rather than additive (linear), as most other basic commodities such as oil, copper, corn, etc., tend to be. Former Google CEO and Chairman Eric Schmidt explains, "Faster airplanes did not help build faster airplanes, but faster computers will help build faster computers." However, the corollary of generative progress is that when the system is starved of the crucial inputs driving the progress (ample, cheap computing power), direct and opportunity costs will also likely feed off each other and grow exponentially.

Our competitors are not standing still. Chief among them, the People’s Republic of China struggles at present to obtain enough compute capability but enjoys ample electricity courtesy of its underutilized coal power plant base and ability to rapidly build wind, solar, and nuclear power plants as well as long-distance power lines. Accordingly, if it continues to make technical progress in chip production, it will be positioned to scale computing power fast and to the detriment of American strategic position. The race is on and the U.S. has no time to waste.

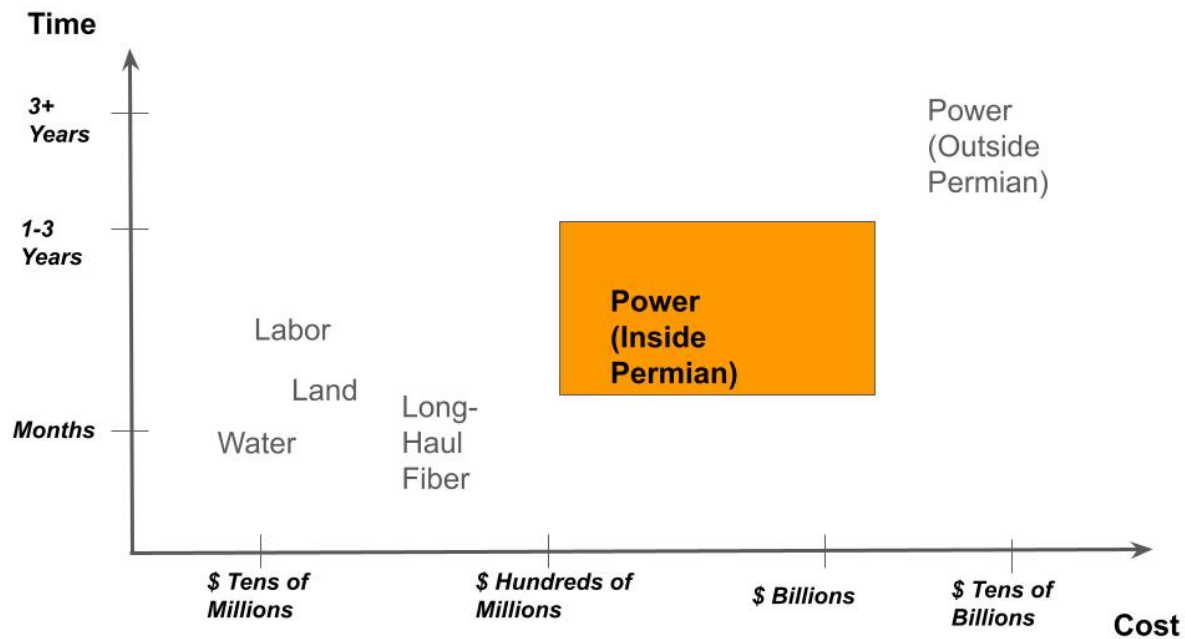
The Permian Basin Can Provide Scalable Digital Infrastructure Solutions

The Permian Basin could bolster American AI in the existential technological competition now underway between China and the U.S. When U.S. leaders formulated the Manhattan Project, they recognized that there was not time to bring world-scale electric power to uranium enrichment and other key nuclear facilities. Instead, they chose to build the facilities where electricity was already available and could be further scaled up in short timeframes. Hence the selection of Hanford, Washington near massive hydropower facilities in the Pacific Northwest and Oak Ridge, Tennessee in the energy-abundant Tennessee Valley Authority powershed.

In today’s battle for global AI supremacy, the Permian Basin offers proven capacity to “build compute near the power.” The Basin offers unique physical space, operational track record, and resource availability to scale up power generation, build high-capacity local power transmission infrastructure, lay additional high bandwidth fiber optic lines, and source water at industrial scale in ways that minimize competition with users such as farms and cities that require freshwater. In a world where speed matters, the ability to rapidly build gas pipelines, power plants, wind and solar farms with batteries, and local transmission lines within the Permian gives the Basin a distinct advantage for AI developers that need to expand computing power as rapidly as possible (Figure 1).

Long haul fiber optic corridors already run through West Texas along Interstate 20 and the adjacent Union Pacific rail line and could be connected to future hyperscale data centers in the Basin via spur lines. The existing long-haul fiber optic “data pipelines” are brownfield routes that can also be expanded faster and likely, at far lower cost, than can pipelines that would take gas to other parts of Texas or the U.S. to power data centers in those locations.⁴

Figure 1 – AI Infrastructure Bottleneck Factors vs. Estimated Cost and Time To Address



Source: Author’s estimates.

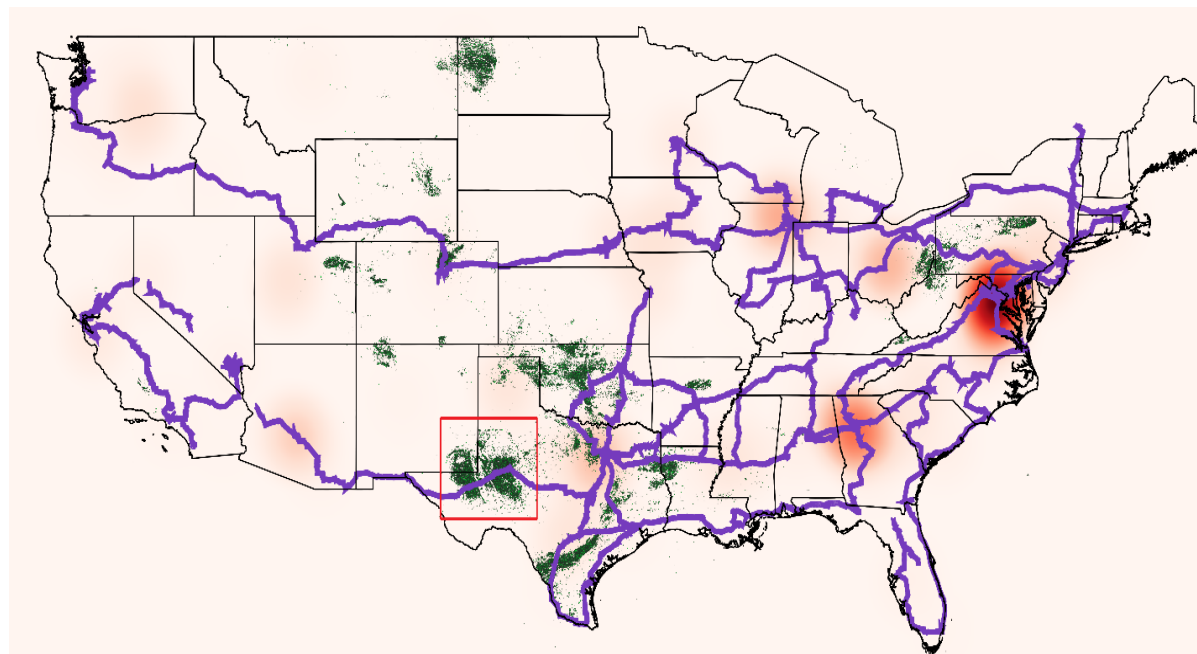
The Basin already boasts top-tier solar, wind, and natural gas resources capable of powering a material portion of the U.S.'s growing data center needs. It offers room to build much more power generation near the source of the lowest cost natural gas in the U.S. It also possesses the physical expanses and distances from population centers to facilitate construction of air-cooled or otherwise low water use small modular nuclear reactors. A more disciplined pace of oil & gas development and reduction in oil & gas worker counts due to upstream consolidation would also augment the digital workforce with fabrication, welding, electricians, and other high-skill tradespeople needed to build and operate digital infrastructure and associated energy assets. Datacenters and the associated local service and supply chains offer a chance to reskill and upskill potentially thousands of workers in a multi-decade industry that, unlike the oil & gas space, is more insulated from commodity price cycles.

Challenges exist, including high summer temperatures, latency requirements, competition for energy (e.g., oilfield electrification, crypto currency mining and potentially, direct air capture of CO2), grid constraints, and scarce freshwater. But the opportunity upside confers urgency upon the search for new operational concepts. Cross-pollinating legacy and emerging energy sources, ample physical space, and non-traditional water resources, including oilfield produced water, can help the Permian Basin become a hyperscale AI powerhouse, particularly for latency-tolerant applications like generative AI.

Thus far, the Permian basin region has virtually no major data center development and no operational hyperscale data centers. New Era Helium intends to build a 250-

megawatt facility near Odessa, Texas but the project is still in the preliminary assessment phase.⁵ As the heatmap below indicates, northern Virginia remains the premier U.S. (and global) datacenter siting location, with smaller concentrations in Ohio, Chicago, Atlanta, Dallas, Phoenix, Salt Lake City, LA, the Bay Area, and Washington state (Figure 2).

Figure 2 – The Permian Has Major Upside for AI Development



Source: Aterio, FracFocus, GADM, Windward, and author’s analysis.

Note: Existing/Under Construction/Announced Data Center Concentrations (Red = most datacenter megawatts in heatmap), selected trunk Fiber Routes (purple), Oil & Gas Wells (dark green).

Powered Land, Gas Competition


One of the key factors that determine the Permian Basin’s AI compute “carrying capacity” will be power. Land is not a constraint. Depending on cooling technologies used and physical spacing of data centers, water could also pose a constraint. In a high water usage case, the largest AI-focused data centers could each potentially require 5-to-6 million gallons of water per day.⁶ That is substantial, but how does it compare in oilfield terms? Six million gallons per day is 143 thousand barrels per day.

To put that in perspective, the author’s past research suggests a 143 thousand bpd user would be on par with the average annual water demand of a 10-rig drilling and completion program.⁷ Geospatial distribution of data centers will be very important, but a move toward air cooling would substantially alleviate water-related bottlenecks and again emphasize power as perhaps the most binding physical constraint other than

semiconductor availability. As one proof of concept example, Microsoft in late 2024 launched an AI-optimized, zero water consumption data center design.⁸

Electricity supplies in the Permian at present come from three core sources: gas for dispatchable power, wind and solar for intermittent power, and grid-scale storage batteries for bridging gaps in renewables generation, particularly during brief periods such as summer sunsets when solar ramps off but overall power loads remain high (Figure 3). Of the Basin’s 30+ GW of power production capacity, only about 5 GW are dispatchable (mostly gas plus the Tolk coal plant on the Basin’s northern edge). This is important for powering up future land because for dependability purposes, datacenter operators will very likely want gas-fired power available. Even if developers are signing power purchase agreements with renewable projects as offsets, unless there are very substantial associated battery assets, they will need dispatchable resources to ensure 24/7 electron availability. These realities have major implications for natural gas use in the Basin, particularly if developers choose to site power generation assets behind the meter to facilitate faster development of projects.

Figure 3 – Permian Basin Installed Electricity Generation Base, MW

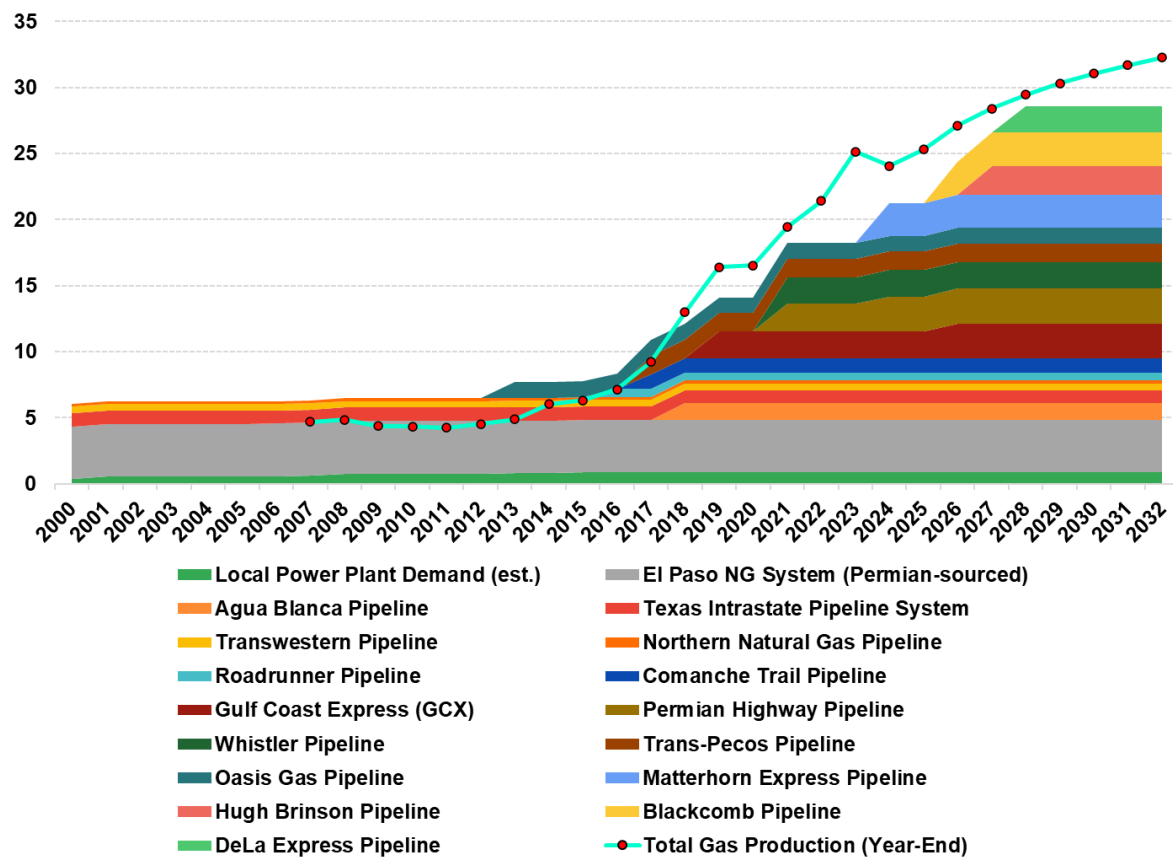


Source: Global Energy Monitor, Cleanview, ERCOT, and author’s analysis.

Gas is attractive for its dependability, scalability, and low local fuel costs. But the Permian is also an increasingly important source of natural gas moved the Gulf Coast by large pipelines and destined for LNG exports. Based on current pipelines in existence and under construction, what might gas availability look like in 3-to-4 years assuming no major supply changes?

With the gas molecules not committed to pipelines, how many megawatts of datacenter capacity could be supported? Let's run some rough numbers. EIA data for gas turbine power plants in the U.S. show a heat rate of 11,010 BTU/kWh in 2023.⁹ Multiplying this by 10³ yields 11 million BTU per megawatt-hour. Assuming 1,055 cubic feet of gas per million BTU, this yields 11,605 ft³ of gas used per MWh of power generated. Assuming capacity utilization of 90%, each MW of gas-fired capacity produces 21.6 MWh of electricity and would consume 0.251 million ft³ of gas. At this ratio, 1,000 MW of gas-fired generation running at 90% utilization would consume 0.251 BCF/d. Accordingly, 1 BCF/d of gas supply could theoretically power about 4 GW of datacenter capacity.¹⁰


Figure 4 – Permian Gas Pipeline Export Projects vs. Local Gas Production, BCF/d



Source: EIA, Energy Transfer, Global Energy Monitor, Kinder Morgan, Inccorrys, Local Media, and author's estimate.

A key question for ensuring power for datacenters hinges on the availability of “surplus” natural gas. By “surplus,” this report refers to natural gas volumes that exceed available pipeline takeaway plus what is plausibly absorbed by existing local gas-fired power plants. The forecasts this author locates suggest that surplus gas volumes will decline between now and 2028 as more pipeline projects come online and will then increase after that, potentially reaching levels that could support as much as 25 GW of additional, local gas-fired power generation capacity in the Permian Basin (Figure 5). That is enough for dozens of hyperscale AI compute facilities.

Figure 5 – Potential Power Generation Using Surplus Gas in the Permian Basin, GW (Assuming 90% Utilization and 11,010 BTU/kWh Heat Rate)



Source: EIA, Global Energy Monitor, and author’s estimate.

Potential Gas Competition and Other Wildcards

Datacenters won't have a monopoly on gas supplies – rather, they will have to compete with LNG exports, petrochemicals, and other gas users. It is important to acknowledge this reality. For perspective, consider what each incremental BCF/d of gas can potentially be used for. Variously applied, that BCF/d could: (1) be liquefied into enough LNG to supply Poland’s full import capacity, (2) supply enough gas-fired power to meet more than 1/4 of the ERCOT Coast Load Zone’s peak usage, (3) power about 3 dozen

world-class AI compute facilities, or (4) supply the gas needs of one of North America’s premier oil refining companies with margin to spare.¹¹

If competitive bidding emerged between gas-centric Permian AI datacenters (using both grid power and behind the meter generation) and other consumers further downstream, the data centers are likely the best positioned to win the bidding contest.¹² As a cost of goods and services sold, energy inputs are relatively lower for an AI-focused data center where operating expenses are likely to be perhaps ¼ of capital expenses amortized over time. For these firms, power cost increases caused by higher gas prices could be more easily absorbed by cloud computing businesses that have the deepest balance sheets in U.S. corporate history and who have in recent years operated with average gross margins in the 30% range.

For the AI datacenter shown in the simple model below, a 10% increase in the price of electricity (a proxy for gas prices) would raise total operating plus capital payback expenses by only 1.7% (Table 1). The relationship for a petrochemical plant or LNG trader would be much closer to 1:1 and they would be more price sensitive.

Table 1 – Simple AI Hyperscale Data Center CAPEX and OPEX Model

Data Center Capacity Unit Cost	\$25	\$Million/MW	https://www.latitudemedia.com/news/microsoft-plans-80b-for-data-centers-as-power-constraints-loom/	
Compute Campus Power	150	MW		
Initial Total CAPEX	\$3,750	Million		
% of CAPEX Attributable to Semis	55%			
Silicon CAPEX	\$2,063	million		
Power cost, \$/MWh	\$100			
Annual Power Expenditures	\$131	million		
Other OPEX (assume 0.5% of initial CAPEX)	\$19	million		
Cost of Capital	5.00%			
Term	5	Years		
Desired Gross Margin	30%	3-yr trailing avg from AWS 10-K. https://d18rn0p25nwr6d.cloudfront.net/CIK-0001018724/e42c2068-bad5-4ab6-ae57-36ff8b2aeffd.pdf		
	Simple Annual CAPEX Payback (\$ Million)	OPEX	Simple Margin	Total Amount to Offset, \$ Million
Year-1	\$476	\$150	\$143	\$769
Year-2	\$476	\$150	\$143	\$769
Year-3	\$476	\$150	\$143	\$769
Year-4	\$476	\$150	\$143	\$769
Year-5	\$476	\$150	\$143	\$769

Source: AWS, Latitude Media, author’s analysis.

If developers use an “all resources on deck” portfolio approach that combined wind, solar, longer-duration batteries, natural gas, and also advanced nuclear reactors, the Permian Basin could potentially host 25-to-50 GW of data center capacity.¹³ Permian oil producer activity is a key wildcard given that the majority of the basin’s gas production

is “associated” and comes as a by-product of oil & liquids output rather than from dedicated gas wells.

If for some reason, oil production peaks or declines sooner than anticipated, or is curtailed due to price responses, the Basin retains great attractiveness for datacenter siting but would be more reliant on renewables and potentially later, nuclear energy. It is less likely, but still plausible, that in a future higher gas price world where carbon sequestration has been proven at scale, coal plants could also be built in the Permian near depleted oil & gas fields and enhanced oil recovery projects that could both house captured carbon dioxide.¹⁴

The anticipated ultimate size of Meta’s data center project in Northeast Louisiana suggests the physical footprint of the total hostable Permian datacenter capacity could be 10 times or more larger than Manhattan.¹⁵ In Texas terms, the ultimate land occupied could be something on the order of the city area of San Antonio or Houston. Solar + battery generation capacity (a current solution) and nuclear reactors (a future solution) paired with data centers can theoretically be built anywhere in and near the Delaware and Midland Basins. In short, a Houston-sized land footprint in a nearly South Dakota-sized basin suggests land is not likely to be a binding constraint on large scale data center development.

Gas power plants would be more closely tied to the availability of gas molecules in the Basin, with potential commensurate needs for gas gathering and transmission infrastructure. Incremental gas collection infrastructure expansions anchored by gas generators paired with data centers can also potentially help reduce/eliminate flaring in the Permian Basin, which in 2025 could amount to 500 million cubic feet/day (enough to support about 3 GW of gas-fired generation running full time).¹⁶

Conclusion

Just as Manhattan Project uranium enrichment facilities succeeded because they were built where the electricity was (or where it could be quickly expanded), so too do AI data centers need to seek areas where silicon can be fed with electrons as quickly as possible. They also need a capable workforce. While tech would be a newcomer to the Permian Basin, the rapid human capital scaleup that occurred during the Shale Boom with tens of thousands of workers migrating to Midland, Odessa, and surrounding communities suggests that when the economic opportunity signal is sufficiently strong, workers will respond.

To that point, data center-related salaries could potentially be competitive with the healthy wages paid throughout the oilfield. As an example of what data center positions could potentially pay, a 2024 local news article discussing xAI’s Memphis data center showed multiple six-figure positions advertised as hiring (Figure 7).

Figure 7 – xAI Memphis Salary Ranges Screenshot, July 2024

Memphis jobs and pay:

- **Foreman**, \$135-165 per hour
- **Lead Fiber Technician**, \$115-135 per hour
- **Network Engineer**, no pay listed
- **Project Manager**, \$100-125 per hour
- **Protective Intelligence Lead**, \$90,000-120,000 annually
- **Security Operations Lead**, \$110,000-140,000 annually
- **Security Technology Lead**, \$110,000-150,000 annually
- **Superintendent**, \$165-240 per hour
- **Fiber Foreman**, \$110,000-\$150,000 annually

Source: WREG News Channel 3. “xAI Now Hiring for Jobs at Memphis Supercomputer Site.” WREG, July 10, 2024. <https://wreg.com/news/local/xai-memphis/xai-now-hiring-for-jobs-at-memphis-supercomputer-site/>.

Data centers are also multi-year assets that would likely have significantly more job stability than the oil & gas industry with its commodity price volatility. AI compute demand appears poised to grow for years to come and once companies deploy multibillion dollar data facilities, they are likely to run them steadily regardless of day to day, month to month economic shifts given the assets’ massive upfront capital cost.

The Permian Basin has great potential to bolster U.S. global AI competitiveness and the recent White House emphasis on gas-rich Pennsylvania as an AI hub suggests the Federal government recognizes the importance of being able to rapidly scale powered land. A confluence of federal, state, and private sector interests makes fertile ground for leveraging the Permian Basin’s advantages as part of an AI Manhattan Project.

Acknowledgement

The author thanks Joseph Webster of the Atlantic Council for his keen insights on the strategic context and framing of the evolving U.S.-China AI competition. He elaborates on many of these concepts in a [recent analysis](#) for War on the Rocks.

Notes

¹ The author is the Water-Energy Thrust Leader at Rice University's Water Institute and holds a membership interest in Cactus Water Services, LLC. This relationship is covered by a Rice University conflict of interest management and monitoring plan.

² "Microsoft earmarks \$80bn for AI data centres in 2025." Total Telecom, January 7, 2025. <https://totaltele.com/microsoft-earmarks-80bn-for-ai-data-centres-in-2025/>.

³ Exxon Mobil Corporation. "Corporate Plan Update and Upstream Spotlight." Investor Relations, December 11, 2024. Accessed August 13, 2025.

https://d1io3yog0oux5.cloudfront.net/_ca7e4e8a5d0aa82a3a315575f0258b12/exxonmobil/db/2261/22346/file/Corporate+Plan+Update+-+FINAL.pdf.

⁴ See, for instance: "BofA Securities 2024 Leveraged Finance Conference," Unifi Investor Presentation, 3-4 December 2024, <https://investor.unifi.com/static-files/30919fea-5545-4bc2-a3e6-d4b584ffa3e9> and Rick Smead, "Natural Gas Pipelines Part 2: Approvals, Financing, and Assessment of Major Projects," RBN Energy, Slide 16, https://rbnenergy.com/sites/default/files/FromGroundUp_Part2.pdf.

⁵ Black, Doug. "New Era Helium Advances 250MW Permian Basin AI Data Center." insideHPC, June 5, 2025. <https://insidehpc.com/2025/06/new-era-helium-advances-250mw-permian-basin-ai-data-center/>.

⁶ Bonk, Lawrence. "Meta Announces Huge New Data Centers, but They Could Gobble Up Millions of Gallons of Water per Day." *Engadget* (via MSN), accessed August 7, 2025.

<https://www.msn.com/en-us/money/companies/meta-announces-huge-new-data-centers-but-they-could-gobble-up-millions-of-gallons-of-water-per-day/ar-AA1Au5m>.

⁷ Gabriel Collins, "How Much Water Does Apache Potentially Need to Develop Alpine High?," Texas Water Intelligence™, Water Note #5, 19 June 2017. [How Much Water Does Apache Potentially Need to Develop Alpine High? – Texas Water Intelligence](https://www.texaswaterintelligence.com/2017/06/19/how-much-water-does-apache-potentially-need-to-develop-alpine-high/)

⁸ [Sustainable by design: Next-generation datacenters consume zero water for cooling | The Microsoft Cloud Blog](https://www.microsoft.com/en-us/blog/2025/07/sustainable-by-design-next-generation-datacenters-consume-zero-water-for-cooling/).

⁹ U.S. Energy Information Administration. "Table 8.2. Average Tested Heat Rates by Prime Mover and Energy Source, 2013–2023." *Electric Power Annual*. Accessed August 13, 2025. https://www.eia.gov/electricity/annual/html/epa_08_02.html.

¹⁰ This report assumes simple cycle turbines which are less efficient, but more rapidly available and deployable for behind the meter power generation. Combined cycle facilities are much more efficient, with heat rates of less than 8 million BTU/MWh rather than the 11 million BTU/MWh of the simple cycle generation. For an example of simple cycle turbines being deployed behind the meter to power a large AI data center, see: Belanger, Ashley. "xAI data center gets air permit to run 15 turbines, but imaging shows 24 on site." *Ars Technica*, July 3, 2025. <https://arstechnica.com/tech-policy/2025/07/xai-gets-an-air-permit-to-power-its-supercomputer-but-pollution-fears-remain/>.

¹¹ LNG data calculated as follows: 1 BCF/d = 0.021 MMT/LNG or 7.665 MMT annualized ([Approximate conversion factors](#)). Poland has 8.3 BCM/284 BCF of LNG intake capacity or about 778 million ft³ per day. ERCOT load zone data obtained from: Gabriel Collins, "Water Load vs. Electricity Load in Texas," Collins Research Portal, 2 December 2024. <https://collinsresearchportal.com/wp->

content/uploads/2025/08/Collins_CES_Water-Load-Analysis_December-2024.pdf. AI compute estimate relies upon the estimates earlier in this paper of what 1 BCF/d could generate in gas-fired power facilities with a 11 MMBTU/MWh heat rate running at 90% utilization. Refinery gas needs estimated based on Valero 2010 public disclosure of actual gas use (<https://www.reuters.com/article/markets/oil/valero-weighs-natural-gas-play-ceo-idUSN29116064/>) compared with estimates from the company's 2024 ESG and 10-K reports, which state energy use of 0.38 MMBTU/bbl of throughput (2,912 kbd in 2024), converting MMBTU to gas use at 1,020 ft³ per MMBTU and then assuming that as a baseline, 50% of energy comes from direct use of gas and other gas use is driven by hydrogen production.

¹² The author notes that this is a simplified discussion of a highly dynamic market that has shown high price responsiveness, is now globalized, and which also features an important presence of Canadian pipeline gas imports into the U.S. as a market backstop. Kenneth B. Medlock III, "Scenarios for Global Natural Gas Markets to 2050: The Dynamics of US LNG Exports, the Deepening Connection Between Oil and Gas Production, and Shifts in Global Demand," Rice University's Baker Institute for Public Policy, March 20, 2025, <https://doi.org/10.25613/ZAET-5W61>. Sustained higher gas prices would also likely stimulate greater use of alternative technologies such as enhanced storage plus renewables and advanced nuclear. The response time for these would be measured in years, whereas the already highly developed North American gas system can respond to supply/demand shifts in days, weeks, and months.

¹³ Gabriel Collins, "Small Modular Reactors for Nuclear Desalination and Cogeneration in the Permian Basin," Rice University's Baker Institute for Public Policy, May 7, 2025, <https://doi.org/10.25613/M0CA-RR71>.

¹⁴ As an historical example of this potential, please see: https://www.gem.wiki/Texas_Clean_Energy_Project.

¹⁵ Abugov, Josie. "Plan to Power Meta's Massive Louisiana Data Center Examined in Key Hearing." NOLA.com, July 16, 2025. https://www.nola.com/news/environment/meta-ai-data-center-louisiana-enviro/article_beb61093-ba7e-4b61-98d9-614f5a1867fe.html.

¹⁶ Talley, Chris. "Permian Basin Pivot: Drilling Slowdown Spurs Operators to Eye Natural Gas." Your Basin, July 16, 2025. <https://www.yourbasin.com/news/permian-basin-pivot-drilling-slowdown-spurs-operators-to-eye-natural-gas/>.