



Clean Energy's AI Moment

How AI's Willingness to Pay can
Reshape Clean Energy Investment



August 2026

Disclaimer

The information provided in this document is for informational purposes only and does not constitute a solicitation, offer, or sale of securities. Neither the investment examples cited nor CREO's mention of examples constitute investment advice or a recommendation to purchase or sell any securities. CREO is not and does not provide services as an investment advisor, investment analyst, broker, dealer, market-maker, investment banker, or underwriter. CREO does not receive any compensation or fee for citing investment examples in this document or any consideration because of any discussion or transaction with respect to any such investments.

No Conflict Statement

The authors report no conflicts of interest related to this study.

Authors

David Siap
CREO Syndicate

Hans Royal
Originator of the Compute Heat Rate (CHR)

Executive Summary

Data centers could be the best offtakers clean energy has ever had, not because they are virtuous, though many hyperscalers carry ambitious clean energy targets, but because of their economics. The economic value of AI is so high that power cost is not material to profitability, especially for higher margin activities like frontier inference, even when factoring in non-electricity costs (e.g., compute hardware, land).

This becomes clear through the lens of two metrics: The Compute Heat Rate (CHR) and the Marginal firm Capacity Cost (MCC), analyzed in the context of the data center heavy PJM Interconnection. The CHR sets the demand ceiling for willingness to pay for power and the MCC assesses the supply side at the system level. What they reveal is that this new class of demand from AI-driven data centers can pay many times the system-level cost of capacity while being highly profitable. The lowest margin compute workload can afford to pay double the cost of the most expensive supply-side resource, and the highest margin workload can pay over 150 times that cost.

Natural gas is a known quantity for developers, with a reported 143 GW of capacity in the pipeline to serve data centers. However, when cost fades as a binding constraint, that naturally elevates other elements of risk, such as duration, firmness/reliability, and fuel exposure. Among those factors, clean resources hold a structural edge and better align incentives with other grid stakeholders. The choice of whether gas or clean resources will power AI deployment at scale is currently live and long-horizon investors are one of the few groups positioned to influence this buildout.

For investors, this maps to six distinct capital allocation postures spanning a spectrum from shortest to longest horizon: demand-side flexibility, behind-the-meter and hybrid generation, storage sited at congestion points, grid-enhancing "picks and shovels," conventional front-of-the-meter generation, and first-of-a-kind advanced generation technologies. Each posture monetizes a different piece of the same underlying dynamic, and the framework does not require picking a single winning technology. Instead, it requires recognizing that cost has stopped being the binding constraint across nearly all of them.

Thus, how the power sector expands to meet demand from AI is a choice, one that does not anchor solely on price. Currently, the data center industry is investing in and scaling up new clean firm technologies while deploying GWs of inefficient gas turbines. Investors are in a rare moment to steward the build out intentionally in a way that respects unpriced externalities and builds a sustainable grid as opposed to one that misprices risk and carries hidden costs.

Table of Contents

Executive Summary	ii
1 Introduction: Two Metrics, One Moment.....	1
2 Methodology	3
2.1 The Compute Heat Rate	3
2.2 The Marginal Firm Capacity Cost	4
2.3 The CHR-MCC Ratio	6
3 PJM Deep Dive: The CHR-MCC Ratio in Practice.....	8
3.1 The PJM Context	8
3.2 The Dominion Hub: Preliminary Evidence	8
3.3 CONE-to-CHR Price Migration	9
3.4 The CHR-MCC Ratio at PJM	9
4 Capital Allocation Implications for Clean Energy Investors	12
4.1 The Investable Ratio	12
4.2 Six Investment Postures	12
4.3 PPA Economics Under the CHR-MCC Ratio	15
5 Conclusion – A Choice The Grid Can Afford To Make	16
5.1 The Clean Energy Thesis, Restated	16
5.2 What to Watch	16
5.3 The Choice Before Us.....	17
References.....	18

List of Figures

Figure 1: MCC curve for Dominion through 2039	6
Figure 2: CHR workload tiers on the MCC curve	9
Figure 3: CHR-MCC ratio by resource	11

List of Tables

Table 1: CHR by workload type, Q1 2026. Gas heat rate benchmark ~\$50/MWh.	3
Table 2: CHR sensitivity to workload composition assumptions.	4
Table 3: Summary of investment postures discussed	14

1 Introduction: Two Metrics, One Moment

A new demand class is reshaping the U.S. power sector. Lawrence Berkeley National Laboratory projects that power demand from AI-driven data centers will more than double from 2023 to 2028.¹ However, supply is struggling to keep up: PJM's capacity auction fell 6,500 MW short; ERCOT's interconnection queue exceeded 238,000 MW; excess capacity on the grid is limited and the thirst for new power is insatiable.² Problems ensue: 14 states have considered moratoriums on new data center construction (including notably Texas in August 2026), community opposition is rising, and central projections have residential electricity prices rising 15 to 40 percent. This leads some experts to discount the current data center pipeline by up to 90%.

Many frame this as a scarcity problem: how does supply catch up with demand? But that framing obscures the more important question. The grid will expand to serve this load, and it will be one of the largest capital deployments in the history of the U.S. power system. The real question is not whether the grid will expand, but which grid gets built. That decision is not made at once but by many actors (developers, utilities, regulators, politicians) running on different clocks, which positions long-horizon capital as one of the few stakeholders positioned to influence the choice at scale and duration.

There are two poles defining this choice, with the likely outcome somewhere between them:

1. One features clean technologies at scale where energy is cheap and the externalized harms of the power sector are a fraction of what they are today. The green shoots of this are the investment in and offtake from the hyperscalers in new clean firm power options such as next-generation geothermal, advanced nuclear, and long duration energy storage (LDES), among others as well as accelerating deployment of renewables and storage.
2. In the other, the energy industry looks much like today, carbon-intensive, resource constrained, and with today's externalities. Industry is now in the early stages of committing to this pathway, as evidenced by the 143 GW of mostly inefficient behind-the-meter (BTM) gas in the current pipeline.³

The Union of Concerned Scientists puts the value of those externalities, and hence the choice, at up to \$13T.⁴ And yet, AI margins are large enough to internalize those externalities without destroying short term profitability, as shown by two metrics.

The first metric focuses on the supply-side cost of new capacity. The Marginal firm Capacity Cost (MCC) framework provides this: a system-level metric that ranks generation resources by their incremental impact on household electricity bills, capturing the full cost of adding one unit of firm peak capacity, including transmission, integration, and reliability costs. MCC shows that the choice of resources used to meet grid expansion has a sharp impact on costs. The lowest cost choices lead to a 5% decrease in ratepayer bills through a period of rapid grid expansion, while a business-as-usual scenario leads to a 27% increase and a ~\$570B increase in capital expenditures through 2035.

The second covers the demand-side price tolerance of the load served. The Compute Heat Rate (CHR) framework⁵ provides this: a metric that quantifies the maximum electricity price an AI data center can rationally sustain before the computation becomes uneconomic. The blended CHR across commercial AI workloads is approximately \$6,350/MWh, or 127 times the gas heat rate benchmark (or the average wholesale electricity price in the market). Even

the lowest-margin category (enterprise contracted) implies a price tolerance of roughly \$1,270/MWh, which is 25 times the gas heat rate.

The combination of CHR and MCC integrates supply and demand signals into a single view. It tells investors which resources pencil for which workloads, at which locations, and by what margin. A large CHR-MCC ratio means the data center's price tolerance vastly exceeds the cost of supplying it: a green light for capital deployment. A small ratio signals risk. A very large ratio shifts priorities yet again, underlining the value of time to power and other strategic benefits.

Under these large ratios, data center developers and investors that set strategic direction are thus free and even incentivized to choose power sources that reduce time to power, avoid fuel-cost uncertainty, and ultimately produce a clean future. The question is, how live is this choice? How big are margins and how often do they occur? What are the risks and opportunities in the gas build out, whether it reverses or endures?

This paper introduces the CHR-MCC ratio as an integrated capital allocation framework. Section 2 presents both methodologies side by side. Section 3 applies the framework to PJM, the largest data center market, where the impacts of these are most visible. Section 4 translates the ratio into actionable capital allocation guidance for clean energy investors.

2 Methodology

2.1 The Compute Heat Rate

The Compute Heat Rate quantifies the maximum electricity price a data center operator can rationally pay before the computation becomes uneconomic. Equation 1 below describes CHR for a given AI workload w :

$$\text{CHR}(w) = (R(w) - C_{\text{non-elec}}(w)) / (1+m) \quad \text{Eq. 1}$$

where $R(w)$ is the revenue per MWh of electricity consumed (derived from API pricing, GPU throughput, and power consumption); $C_{\text{non-elec}}(w)$ is the non-electricity operating cost per MWh (GPU amortization, facility costs, networking, cooling); and m is the required margin (30% is common).

$R(w)$ comes from a four-link conversion chain: the CHR framework converts published API prices (\$/million tokens) to throughput (tokens/second from MLPerf benchmarks), then to power consumption (kW per server from hardware specifications), yielding revenue per MWh. The derivation uses only publicly available pricing and hardware data. Table 1 reports these values per watt of annual capacity, converted at a 100% load factor.

The resulting CHR values span three orders of magnitude across workload types:

Table 1: CHR by workload type, Q1 2026. Gas heat rate benchmark ~\$50/MWh.

Workload Type ⁶	R(w) \$/W/year	C_non-elec(w) \$/W/year	CHR \$/W/year	vs Gas HR
Frontier Inference	\$648	\$37	\$470	~1,070x
Mid-Tier Inference	\$130	\$37	\$71	~162x
Enterprise Agentic AI	\$131	\$39	\$71	~162x
Enterprise Contracted	\$52	\$37	\$11	~25x
Commodity Inference	\$16	\$33	\$7*	~16x
Blended (Q1 2026)	\$110	\$38	\$56	~127x

*Dollar values converted from the Royal (2026) \$/MWh basis to annual capacity terms at 8,760 hours per year (100% load factor); the gas heat rate multiples are unchanged by the unit conversion. 30% required margin applied. *Effective CHR reflects portfolio-level cross-subsidization. Source: Royal (2026), [SSRN 6322318](#).*

Sensitivity to workload composition

The blended CHR depends on consumption-weighted workload shares estimated from hyperscaler earnings disclosures and industry analyst reports, not directly observed from proprietary operator data.⁷

Table 2: CHR sensitivity to workload composition assumptions.

Scenario	Blended CHR		vs Gas HR	Assumption
	\$/W/year	\$/MWH		
Full blend	\$56	\$6,350	~127x	All five tiers at estimated weights
Conservative	\$33	~\$3,800	~76x	Agentic tier eliminated; weight redistributed to enterprise contracted
Floor	\$11	\$1,270	~25x	Only lowest tier

The conclusion is the same across all three scenarios: even at the floor case, the CHR exceeds the MCC of every generation resource on the supply curve by a wide margin. A clean energy investment thesis does not depend on the 127x figure being precisely correct. The ratio is so large it survives even the most aggressive haircut.

The CHR values referenced in this paper are from the Q1 2026 CHR index. CHR is a living metric, constantly moving with new model announcements, API price changes, user workload mix, and many other factors. As a preview, the next CHR index likely will show lower values across several tiers, especially in the mid- and low-CHR tiers¹. Multiple forces are at work: commoditization of inference with open weights and smaller models absorbing large swaths of routine work; downward price pressure as providers compete; and stratification between the ultra-frontier and commodity tiers as the market continues to mature. None of these movements change the thesis of this paper, as the gap between AI electricity price tolerance and traditional loads is measured in orders of magnitude. Quarterly movements in CHR do not alter the conclusions herein.

2.2 The Marginal Firm Capacity Cost

The MCC measures the incremental effect of adding one unit of firm peak capacity from a given resource on the average household's annual electricity bill. The metric reflects the fact that peak demand determines grid build out, and capacity shortfalls typically refer to specific stressed periods. Where the CHR is a demand-side ceiling, the MCC is the supply-side floor: the all-in system cost of actually delivering reliable power. It is a system-level metric that, unlike LCOE, includes capital cost, fixed and variable operating cost, the resource's reliability contribution, the demand-reduction value of efficiency and flexibility, and avoided or deferred transmission and distribution (T&D) investment into a single figure. While the original metric focused on household bills, for the purposes of this paper the framework has been adapted to report capacity costs rather than bill impacts while maintaining some of the ratepayer focus.

Each of 32 resources appears as a bar. Its width is the new nominal capacity potential in the planning horizon, drawn from published central estimates; its height is the ratepayer cost per unit of firm capacity annually. Resources rank from lowest to highest to form a marginal cost curve, directly analogous to the marginal abatement cost curve of McKinsey (2007) but

¹ Quarterly CHR index values, tier definitions, and update history are published at www.computeheatrate.com

ordered on the cost of firm peak capacity rather than carbon. The height of each bar begins with the annualized, reliability adjusted cost of capacity as in Equation 2:

$$\text{MCC_ARR} = (\text{CapEx} \times \text{FCR} + \text{FOM} + \text{Fuel}) / \text{ELCC} \quad \text{Eq. 2}$$

CapEx is overnight installed cost per kilowatt. The fixed charge rate (FCR) annualizes capital: in its regulated form it is the capital recovery factor, $\text{CRF} = r(1+r)^n / [(1+r)^n - 1]$ for cost of capital r and asset life n , grossed up for the depreciation tax shield and for ongoing property-tax and insurance charges. The methodology adds fixed O&M and fuel costs where applicable after annualization.

Effective load-carrying capability (ELCC) reflects what portion of nameplate capacity is effectively firm capacity. Because ELCC degrades as penetration of like resources increases, the curve is penetration-dependent: each incremental gigawatt of solar or short-duration storage is worth less than the last. This is what places resources that look cheap on an energy basis (e.g., rooftop solar) high on the capacity curve due to their limited marginal direct contribution during peak times.


The capacity cost folds into the utility's annual revenue requirement (ARR) and nets against two credits to arrive at a final annual per firm GW cost as in Equation 3:

$$\text{MCC} = (\text{MCC_ARR} - \text{Credit_demand-reduction} - \text{Credit_T\&D}) \quad \text{Eq. 3}$$

The demand-reduction credit captures the impact of avoided energy use via efficiency and demand response, which is what drives those resources to negative cost at the far left of the curve. The avoided-T&D credit nets out transmission and distribution investment deferred or avoided by resources that shave peak or sit close to load. The rate base remains constant as a conservative simplification, omitting the bill relief that comes from spreading fixed costs across rising sales.

Figure 1 shows the resulting curve for Dominion territory. It takes local sources such as Dominion's integrated resource plan (IRP) into account, which runs through 2039, as well as central literature estimates and other sources. The analysis includes all 32 resources as in the original paper, however not all resources are likely to have significant new capacity during this period in this territory (e.g., new hydro), so the width of their bar is set to zero. Finally, while Commonwealth Fusion Systems' 400 MWe ARC plant targets COD in this territory in the 2030s, public cost data is not available at time of publication, preventing inclusion here.

Figure 1: MCC curve for Dominion through 2039



Like the national results,⁸ demand flexibility results in positive value to ratepayers when accounting for reduced demand and offset grid infrastructure build out. Batteries and utility scale solar dominate nominal capacity additions, again mirroring the national case. However, since the ELCC for solar in Dominion territory by 2039 is substantially lower than the projected national value by 2035, projected costs for both utility-scale and rooftop solar are higher. Still, solar remains an incredibly cheap source of energy, which places it central to the energy transition. Fossils provide somewhat cheap capacity but carry externalized carbon costs and are subject to volatile fuel pricing. Finally, offshore wind and SMRs are relatively costly sources of capacity expected to add over 1GW of capacity each in the 2030s.⁹

2.3 The CHR-MCC Ratio

The CHR-MCC ratio for workload w and resource r at location l is:

$$\text{Ratio}(w, r, l) = \text{CHR}(w) / \text{MCC}(r, l) \quad \text{Eq. 4}$$

A ratio over one indicates the data center's electricity price tolerance exceeds the full system cost of a generation resource. A ratio between one and zero indicates the resource does not pencil for that workload class. The magnitude measures the margin of safety: a wide positive ratio means the economics are robust to cost overruns, regulatory changes, and market shifts.

Two caveats keep the ratio honest. CHR is a price ceiling, not a clearing price: wholesale prices will generally not rise to meet it, so the ratio measures structural headroom—how far demand sits above the cost of supply—rather than a return any investor captures. And the signal does not hinge on the blended figure; even the floor-case workload (enterprise contracted, CHR ~\$11/W/year) sits well above the full system cost of firm generation, leaving the ratio firmly positive. The blended \$56/W/year is upside, not the anchor. The two values also sit on different balance sheets, where CHR is the buyer's private willingness to

pay and MCC is the all-in system cost of supply. The ratio measures the value of the load against the true cost of serving it, not against any single party's bill.

The ratio is location-specific because MCC varies by geography, reflecting the grid realities of the local rate base in the local utility's fixed charge cost (and embedded cost of capital, r), the local ELCC which is dependent on the local grid and renewable resource, as well as the potential capacity available in the timeframe, which varies greatly based on the local resource. It is workload-specific because CHR varies by AI task. And it is time-varying because both CHR (which tracks AI pricing and hardware evolution) and MCC (which tracks construction costs, fuel prices, and policy) are dynamic quantities.

3 PJM Deep Dive: The CHR-MCC Ratio in Practice

PJM is the natural laboratory for the CHR-MCC framework. It is the world's largest wholesale electricity market and also contains the world's largest data center cluster (Northern Virginia, Dominion zone). And, as of May 2026, it is the only market operator to have cited both the demand-side and supply-side dimensions of the AI electricity pricing problem in the same analysis.

3.1 The PJM Context

PJM's May 2026 white paper on market design reform explicitly cited the CHR framework (pp. 49-50) in the context of scarcity pricing reform.¹⁰ The same paper referenced the cost of new generation capacity, drawing on GridLab and Halcyon data to estimate the cost of gas-fired generation. PJM is confronting the reality that its capacity market must price the interaction between supply-side costs and demand-side price tolerance, not either in isolation.

The numbers illustrate why. PJM attributes approximately 5,100 MW of its 5,250 MW increase in forecast peak load to data center demand. Three consecutive capacity auctions cleared at record prices. The most recent auction fell 6,500 MW short. New gas plant costs have doubled to \$2,100-\$2,300/kW. PJM is simultaneously facing the most price-tolerant demand class in electricity market history and the most expensive supply environment in decades.

3.2 The Dominion Hub: Preliminary Evidence

The PJM Dominion Hub provides the most concentrated natural experiment for CHR-MCC interaction. Northern Virginia hosts the world's densest data center cluster. Dominion Energy Virginia's metered billing demand data¹¹ provides direct measurement of data center load at this settlement point.

Preliminary regression analysis at the Dominion Hub shows evidence consistent with the CHR thesis: incremental data center load is associated with meaningful increases in wholesale electricity prices. The estimated per-GW translation suggests a cumulative impact in the range of \$8-14/MWh, depending on specification and time period. This preliminary finding is directionally consistent with independent research on Virginia wholesale markets. Mamkhezri, Sun, and Yang (2025) found a \$2.49/MWh congestion and loss increase attributable to data center growth in Virginia between 2015 and 2024, with each additional data center adding approximately \$0.81/MWh.¹² Still, analysts debate both the magnitude and the direction of data centers' effect on wholesale and retail prices.

Any such price effect would likely persist across all hours rather than concentrate in peak periods, consistent with the flat 24/7 load profile of data center operations. This distinguishes the CHR price transmission mechanism from traditional demand-driven price spikes, which cluster in afternoon and evening peak hours.

3.3 CONE-to-CHR Price Migration

The traditional wholesale price formation model assumes that prices oscillate around the cost of new entry (CONE), currently \$80-\$130/MWh for gas-fired generation in PJM. When prices exceed CONE, industry builds new generation, adding supply, and moderating prices. When prices fall below CONE, no new generation is economic, and retirements tighten supply. This self-correcting mechanism has governed wholesale markets for decades.


CHR breaks this mechanism. A demand class with a \$6,350/MWh price tolerance does not curtail when prices exceed CONE. It does not reduce consumption at \$200/MWh. It does not exit the market at \$500/MWh. The SLA Dual-Lock explains why: data center operators face both an economic lock (the CHR ceiling is far above any plausible wholesale price) and a contractual lock (99.95%+ uptime SLAs with liquidated damages make voluntary curtailment economically irrational).

The result is a price migration from CONE upwards. The speed of this migration depends on data center penetration. At low penetration, the traditional self-correcting mechanism still dominates. As penetration grows, the marginal buyer at the clearing price is increasingly a data center, and the price floor is likely to ratchet up. This is where the traditional model breaks: the load arriving at the margin is large, and its consumption does not respond to grid price signals. Once the marginal buyer stops responding to price, price stops governing the buildout, and the decision of which grid to build can rationally rest on grounds other than cost.

3.4 The CHR-MCC Ratio at PJM

Figure 2 overlays the CHR workload tiers on the MCC resource cost curve for Dominion, adapting both frameworks to report the cost of firm capacity per year. This compares the hyperscaler willingness to pay for capacity versus the ratepayer impacts for a given resource under a static rate base assumption.

Figure 2: CHR workload tiers on the MCC curve



Overlaying the CHR values onto the MCC curve visually shows the dynamic that is threatening to distort the power system as AI scales. The economic value of AI is so high that power cost is not material to profitability, especially for higher margin activities like frontier inference. This holds even when factoring in non-electricity costs (e.g., compute hardware, land) as in C_non-elec(w) in Table 1. That said, there is still some nuance to these results.

There is a material difference between a 2x ratio, as exists between the top of the MCC curve and commodity inference, and a 23x ratio for agentic AI or the 150x ratio for frontier inference. At 2x, a project is still vulnerable to reduced demand or margin compression. At 23x, cost is almost irrelevant, thus power quality, strategic benefits, and time to power increasingly drive procurement decisions. At 150x, the economic value is so high that time to power dominates the selection criteria.

From the CHR perspective, the figure reveals a structural asymmetry. Even the lowest-margin AI workload can absorb resources costing 16x the gas heat rate and 2x the highest MCC resource. The blended CHR (\$56/W/year) is so far above the MCC of any generation resource that the question is not whether the data center can afford the electricity; it is which resource combination delivers the lowest-cost, lowest-risk supply at that specific location.

This reframes the clean energy thesis, and the Fuel-Side Corollary¹³ is what separates gas from clean inside the ratio. On a static MCC curve, gas and clean resources both sit deep inside the CHR envelope, so the ratio looks large for both. It is not equally durable.

Gas carries hidden exposure. Data center load is flat, 24/7, with no seasonal relief, so it raises baseload gas demand. Higher gas demand raises gas prices, which raises the gas plant's operating cost, increasing its MCC. The same demand growth that creates the CHR also feeds gas markets and narrows the gas ratio from below. A gas position at a high-CHR PJM hub is long the very fuel-price risk the buildout creates. Stated as an outcome: gas capacity carries an embedded fuel-risk premium, and that premium is not static; every increment of data center load deepens the fuel exposure of every gas plant already on the system, so the exposure compounds with the very buildout that makes the gas investment look attractive.

Clean resources (solar, wind, nuclear, geothermal, storage) have zero marginal fuel cost, so their MCC does not move with commodity markets. The clean ratio is structurally more stable and widens as the Demand Composition Ratchet intensifies,¹⁴ while the gas ratio compresses. In PJM, where gas typically sets the marginal price, that is the decisive split: the question is not whether a resource pencils today, but whose ratio survives the demand growth that makes the thesis work.

Figure 3: CHR-MCC ratio by resource

Ratio of MCC to CHR-WTP	Workload type	Frontier Inference	Mid-Tier Inference	Enterprise Agentic AI	Enterprise Contracted	Commodity Inference
Resource Type	\$/Wfirm/year	470	\$71	\$71	\$11	\$7
DR - T&D elimination	\$(0.18)	-2620	-396	-395	-62	-39
Efficiency - Motors and process equipment	\$(0.15)	-3159	-478	-476	-75	-47
Efficiency - Lighting and controls	\$(0.12)	-4074	-617	-614	-96	-61
Efficiency - HVAC	\$(0.07)	-6458	-977	-973	-153	-96
DR - T&D deferral	\$(0.04)	-12136	-1837	-1828	-287	-181
DR - No T&D impact	\$0.06	8363	1269	1263	198	125
Gas Peaker - Extended retirement	\$0.09	5262	796	793	125	78
Hydro - Extended retirement	\$0.10	4598	696	692	109	69
BESS - Utility scale - curtailed renewables	\$0.15	3147	476	474	74	47
BESS - Utility scale - off peak	\$0.19	2475	375	373	59	37
Gas Peaker - New	\$0.20	2321	351	350	55	35
Coal - Extended retirement	\$0.21	2293	347	345	54	34
Pumped hydro	\$0.22	2169	328	327	51	32
BESS - Utility scale - balancing	\$0.24	1970	298	297	47	29
CCGT	\$0.33	1419	215	214	34	21
BESS - BTM - onsite solar	\$0.37	1287	195	194	30	19
Nuclear - Extended retirements / restarts	\$0.39	1212	183	183	29	18
Nuclear - Upgrades/expansions	\$0.46	1023	155	154	24	15
BESS - BTM - grid	\$0.49	965	146	145	23	14
Geothermal (conventional)	\$0.52	910	138	137	22	14
Utility solar	\$0.55	860	130	129	20	13
Geothermal (next generation)	\$0.57	828	125	125	20	12
Other LDES	\$0.66	709	107	107	17	11
Hydro - New	\$0.68	693	105	104	16	10
Onshore wind	\$0.71	660	100	99	16	10
Offshore wind	\$0.91	518	78	78	12	8
Hydrogen	\$0.94	502	76	76	12	7
Nuclear - New	\$0.97	484	73	73	11	7
Rooftop solar - comm	\$1.13	418	63	63	10	6
Nuclear - SMR	\$1.21	389	59	59	9	6
Efficiency - With major retrofit	\$1.38	341	52	51	8	5
Rooftop solar - resi	\$3.09	152	23	23	4	2

Demand side flexibility – Maximum impact, challenging reliability

Very high multiples 100x+ – Focus on squarely on time to power

High multiples 10-100x – Select for quality, time to power, and strategic benefit

Low multiples 1-10x – Vulnerable to margin compression

Because the ratio clears so widely, its value is less a pass/fail screen than a map of what governs selection once cost stops binding. Figure 3 organizes this into four bands. Below 1, the resource does not pencil. Between 1 and 10x, cost still binds: a project clears narrowly, and a repricing of the workload or a construction overrun can make it uneconomic, so this is the only band where least-cost selection governs and where committed long-lived capital is most exposed. Between 10 and 100x, cost stops binding and selection turns on time to power and power quality: firmness, duration, reliability, and fuel exposure, the dimensions where clean firm generation and storage generally win. Above 100x, value so dwarfs cost that the ratio is no longer a screening tool; every resource clears, execution and access are the whole game, and every day of delay represents millions in lost revenue, which is also where community opposition bites hardest and pushes projects behind the meter or toward generation with higher community acceptance.

4 Capital Allocation Implications for Clean Energy Investors

The CHR-MCC ratio is a decision tool. This section translates the framework into the capital allocation logic that CREO members and clean energy investors can use.

4.1 The Investable Ratio

For any candidate generation resource at a specific location, the CHR-MCC ratio supports the investment thesis. If the ratio is greater than one, the resource is likely viable for that workload class. Since the ratio is large, the economics are resilient to execution risk, cost overruns, and regulatory uncertainty.

Renewables, storage, nuclear, and geothermal are all evaluated against a conventional industrial buyer whose price tolerance is \$50-80/MWh. Against that buyer, the cost competitiveness of clean resources varies by geography, technology, and policy environment, and the economics can be marginal. Against AI data centers, many clean energy resources on the MCC curve may be deeply investable. The question shifts from "can we compete with gas" to "which clean resource best serves this specific location and workload mix."

One implication converts the Fuel-Side Corollary from mechanism into strategy. Fuel risk is priced project by project: a single gas pro forma runs today's forward curve and hedges near-term volatility, but no individual project prices the cumulative effect of an enormous systemic buildout on the gas demand it shares with every other plant. That is a system-level exposure, hard to see at project scale, and therefore largely absent from the underwriting. For investors, this makes clean energy a no-regrets directional position. A clean project underwritten on today's assumptions has no fuel input to be wrong about. If the Fuel-Side Corollary never materializes, the underwriting holds; if it does, clean assets outperform a gas fleet financed by counterparties who never priced it. The asymmetry costs nothing to hold, and it is available precisely to allocators capable of thinking at the system level rather than the project level, the position long-horizon investors occupy.

4.2 Six Investment Postures

The CHR-MCC ratio maps to six distinct capital allocation strategies, differentiated by how the investor positions relative to the ratio and organized by time-horizon from shortest to longest.

Long demand-side flexibility (shortest horizon)

The framework places demand flexibility at the far left of the curve, cheaper than any generation resource, making it the first place capital should look for the most economic source of headroom. The framework credits demand flexibility with the ratepayer benefits it produces, so it resolves to a negative cost and an undefined ratio. For clarity, this is not the internal load-shifting hyperscalers like Google pursue within their own campuses; it is grid-side headroom produced through vehicles like Google's Energy Impact Funds and its Voltus¹⁵ agreement, and of the kind Rewiring America¹⁶ has proposed, freeing existing capacity rather than building new. The posture is not building a plant but financing and

assembling aggregations of demand and earning on the gap between what curtailment is worth to a capacity-short grid and what it costs to pay participants. It inherently features low capital intensity, can be faster to deploy than generation, and is uniquely additive on community goodwill, since it lowers bills directly. Traditionally, players in this sector have struggled with unit economics¹⁷ and high customer acquisition costs. Technology can address this and in a territory like Dominion facing a tight 2039 reserve margin, it is the cheapest megawatt available. Since much of its value is in avoided costs, it avoids stranded asset risk but also does not produce durable assets with regular returns like many other postures, placing it earliest on the time-horizon axis.

Long behind-the-meter generation and hybrid capacity

If the CHR-MCC ratio is large enough, it becomes rational for the data center operator to own the generation as essential infrastructure rather than procure it through the grid. As stated in the introduction, this is already happening; 143 GW of BTM gas is in the pipeline, and some sources suggest that BTM generation will power 25-35% of AI data centers, either permanently or as bridge power. Referencing **Error! Reference source not found.**, the CHR-MCC ratio¹⁸ for new gas peakers is 35 to 2,321, and for CCGT it is 21 to 1,419.¹⁹ That means it is economic under virtually all future conditions to power AI with gas if externalities are unpriced. It is also economic to power them this way even if regulators eventually force them to retire well before the end of the useful life of the gas plants.

Community pushback is intensifying and advocates have sued data center operators for circumventing local pollution laws.²⁰ Currently, the U.S. Department of Justice has intervened on behalf of the data center industry. However, if that changes in the future, it remains highly economic to retire these already inefficient BTM gas plants in favor of less controversial options. Clean firm power providers may benefit as price makers capturing the majority of the CHR-MCC spread in this case as the stranded assets become not the gas plants but the compute hardware that will need power. If the BTM gas build out intensifies unabated, this increases the value of point-source carbon capture, a technology historically gated by power-sector economics (completely upended by data centers). Finally, since the BTM build out is mostly simple cycle plants, it appears unlikely that these will interconnect to the grid in the future, as the more efficient CCGT has dominated additions over the last 10 years.

Bring your own capacity (BYOC) frameworks, pairing BTM with grid resources, are gaining popularity as they offer a blend of time to power, lower costs, and increasingly lower emissions. The same risks as the fully BTM case apply, but to a lesser extent, at the cost of increased complexity due to the necessary relationship with the local utility.

Long storage

Data center clusters create localized congestion because a small number of settlement points concentrate the load while generation distributes across the system. An investor who can identify hubs where the CHR-MCC ratio is wide and congestion is increasing can capture value by building generation or storage at the congestion point, effectively arbitraging the location-specific ratio. This is the strategy most naturally aligned with clean energy: battery storage, in particular, captures congestion value by performing energy arbitrage or grid services, providing reliability and firmness demanded by high-CHR loads. In the longer term, LDES provides seasonal shifting necessary to power data centers through winter peaks.

Long/Short picks and shovels

One level back from generation: grid-enhancing technologies, advanced conductors, transformers, software that shortens the timeline to COD, and clean energy developers. This posture monetizes the buildout regardless of which resource wins. This stance is more diversified and thus lower beta to the resource-selection question, which is a desirable feature for some allocators. In the short term, these technologies offer obvious value, and many investors are positioning themselves long. However, over an extended time horizon, the picture muddies as excess generation, especially BTM, lowers the need for many of these solutions if utility incentives remain unchanged.

Long Front-of-the-meter (FTM) generation

The conventional play: build generation and sell into the wholesale market or via PPA. This is the highest-volume posture here as it does not require a bilateral agreement with a hyperscaler and represents the majority of power build out to meet data center demand. As shown in **Error! Reference source not found.**, the CHR-MCC ratio is strongly positive for all resource-workload combinations, and all resources are highly positive for the typical workload. Figure 2 shows the potential new capacity in Dominion territory. Together with the ratio, an investor could assess the value at stake for the buildout of each resource and prioritize accordingly.

Long advanced generation technologies (longest horizon)

Frontier inference may be a great pair with frontier clean generation. Hyperscalers are funding advanced nuclear/SMRs, next-generation geothermal, long-duration storage, and fusion directly, taking equity and first-of-a-kind offtake risk precisely because as buyers they prioritize the attributes clean firm power provides. That willingness to absorb the first-of-a-kind premium is a structural tailwind to cross the valley of death until the technology is financeable. This is the highest-beta expression of the thesis and carries real technology and execution risk, but success stories like Fervo Energy highlight the upside.

Table 3 summarizes these investment postures in a more digestible format.

Table 3: Summary of investment postures discussed

Posture	Time Horizon	Capital Intensity	Key Risks
Long Demand Flexibility	Shortest	Low	Execution, unit economics/customer acquisition cost, resource availability
Long BTM and Hybrid/BYOC	Short	High	Regulatory reversal, community opposition, stranded asset, demand fluctuation
Long Storage	Medium	High	Durability of the local premiums, financing
Long/Short picks and shovels	Medium/ Long	Low	Demand cannibalization from BTM build out, incentive mismatch/customer concentration
Long FTM generation	Long	High	Least differentiated posture, potentially merchant risk
Long advanced generation technology	Longest	Medium	Technical/execution risk, FOAK financing risk

4.3 PPA Economics Under the CHR-MCC Ratio

It is worth being precise about what the ratio does and does not promise. No developer captures \$6,350/MWh. A developer signs a power purchase agreement at \$80, \$100, or \$120/MWh, whatever the negotiation yields, and lives on that revenue for 10 to 20 years. The CHR-MCC ratio is not a revenue forecast; it is an underwriting tool. This signals to a clean energy investor that the parties will sign the contract on good terms, the counterparty will not walk away, and a favorable environment awaits re-contracting upon expiration. Each of those is worth money, and each follows from the width of the ratio.

The first effect is on execution. A buyer with thousands of dollars per megawatt-hour of headroom does not spend 18 months negotiating a \$5/MWh concession on a \$90 PPA. The deal closes, and it closes faster and more reliably than a comparable industrial offtake. For a developer carrying interconnection deposits and construction financing while the meter runs, that speed to a signed contract is not a soft benefit; it is a direct reduction in the cost and risk of getting the project built.

The second effect is on price. Because the buyer's tolerance runs from roughly \$1,270/MWh at the floor to \$6,350/MWh blended, a PPA struck at \$100/MWh is a rounding error against the value the electricity supports. The developer can therefore price at a genuine premium to what a conventional industrial offtaker would ever accept, widening the margin between the PPA price and the project's MCC. Wider margins make the project more bankable, easier to finance, and cheaper to capitalize.

The third effect is on durability throughout the life of the contract. A conventional industrial offtaker may walk away if wholesale prices fall below the contracted rate, a well-known risk in the renewable PPA market. A hyperscaler will not. Electricity is one to two percent of the revenue stack for high-CHR workloads, and the operational cost of unwinding a power arrangement, requalifying interconnection, and re-architecting redundancy dwarfs any savings a lower price could offer. The SLA Dual-Lock reinforces the point: uptime guarantees above 99.95%, backed by liquidated damages, make a voluntary supply disruption economically irrational.

The fourth effect is the merchant tail. A PPA signed today expires in year 15, and the question is what the asset is worth after it. Where the hub carries heavy data center concentration, the demand class is still there, because the CHR framework describes a structural feature of the load rather than a cyclical one. Re-contracting into that environment would be quite favorable, and lenders and equity investors capitalize that residual value into the project today as shown by activity in PJM. A CHR-heavy hub therefore carries a more valuable merchant tail than a hub whose offtaker might relocate or fail within 15 years.

These effects compound in the capital stack. Faster execution, a pricing premium, a counterparty that will not walk, and a favorable merchant tail together let lenders underwrite at higher advance rates and let equity accept lower returns. Tax equity and institutional capital that would never participate in a conventional industrial PPA can participate here, because the offtake risk is, by the numbers, de minimis.

The conclusion is straightforward. Data centers are the best offtakers clean energy has ever had, not because they are virtuous, though many hyperscalers carry ambitious clean energy targets, but because their economics make the PPA structurally unbreakable. The CHR-MCC ratio does not merely identify where to build. It transforms financing of the build.

5 Conclusion – A Choice The Grid Can Afford To Make

5.1 The Clean Energy Thesis, Restated

The conventional clean energy investment thesis is cost-competitive generation replacing fossil fuels. The CHR-MCC ratio reveals a second, potentially larger thesis: clean energy as the preferred supply for a demand class that is so powerful that it must rethink electricity procurement. Instead of an undifferentiated commodity where cost is king, these buyers value other attributes such as time to power, reliability, and zero fuel cost exposure above all else.

The data center buyer is not choosing between gas and solar on an LCOE basis but instead selects resources with zero/low fuel cost exposure (solar, wind, nuclear, geothermal). When price tolerance is 127x the gas heat rate, the risk to hedge against is not the cost of electricity; it is the volatility and availability of electricity. Clean energy resources, particularly when paired with storage, minimize exactly those risks.

For CREO members, the implication is that the CHR-MCC framework reframes clean energy from a cost story to a risk story. The buyer's price tolerance removes cost as the binding constraint. What remains is reliability, duration, and fuel exposure, and on all three dimensions, clean energy has a structural advantage.

5.2 What to Watch

The CHR-MCC ratio is dynamic. Several observable signals will determine whether it widens or narrows in the coming quarters.

AI revenue per MWh: if API pricing compresses faster than efficiency gains offset it, CHR declines. If total compute demand grows faster than per-unit cost declines, CHR increases. The CHR index tracks these dynamic trends over time.

Construction costs: rising MCC narrows the ratio from the supply side. Current trend: gas plant costs doubled; renewables declining or stable. Net: ratio is widening for clean energy resources relative to gas.

Data center penetration: The ratio accelerates as penetration grows. PJM is the leading indicator; ERCOT is the fastest growing. MISO is emerging, too.

Policy interventions: moratoriums, cost socialization rules (Texas Governor June 2026), and FERC co-location proceedings all affect how the ratio transmits. These are speed variables, not magnitude variables: they affect how fast the ratio widens, not whether it does.

BTM acceleration: the rate at which hyperscalers shift to self-build generation is a direct measure of how operators are responding to the ratio. Meta \$11B, Google 2.7 GW, Applied Digital integrated campuses. This trend validates the ratio but also shifts where it materializes (from wholesale markets to bilateral agreements and asset ownership).

5.3 The Choice Before Us

This paper opened with two futures. In one, clean, abundant power meets the AI buildout, and the grid emerges more affordable and less carbon-intensive than it began. In the other, the 143 GW of inefficient BTM gas already planned to power data centers are the down payment on a much larger build-out, and the externalities that follow.

The two metrics together establish that the choice between these two futures is live. The CHR shows the demand arriving on the grid can bear power costs one to two orders of magnitude above what any resource on the MCC curve costs to build; even the lowest-value workload, at the most aggressive haircut, clears almost every resource, and the average resource clears it by more than an order of magnitude. Cost has stopped being the binding constraint on the largest power buildout in a generation. The margins are large, they hold across every workload tier, and they are more durable for clean resources than for the fossil capacity they would replace.

When cost stops binding, selection falls to what remains: firmness, duration, fuel exposure, carbon, community acceptance, and speed to power. On each, clean power holds a structural advantage and should resonate with a buyer indifferent to price but sensitive to risk and time, as many are at this moment. The market's logic, not policy mandate or corporate virtue, is what makes the better of the two futures the rational choice. The choice is live and it is affordable.

What remains is execution, and stated plainly, the tension in this paper's own framework is thus: Every actor in the buildout runs on a different clock. The data center developer's clock is the time to commercial operation date as speed to power dominates. The utility's clock is the rate case and the regulator's clock is the docket. The politician's clock is the election. Attempting to persuade developers to slow down misreads the problem; everyone's incentives are functioning as intended.

Long-horizon private capital is suited to take a longer view, since it sits on both sides of the trade, in the capital stacks of the data center operators and in the generation assets that will power them. These investors see the whole board where the grid operator sees only the front of the meter. Their allocation decisions are the steering mechanism itself: the cumulative set of pro formas funded collectively determines which grid gets built, and having a 20-year horizon helps tremendously.

None of this asks investors to be virtuous. The paper's argument has been economics from the first page, and it stays economics here: Section 4.1 showed that at this horizon, the long-term-optimal allocation and the return-maximizing allocation are the same trade. Clean assets carry no unpriced system-level fuel risk. Gas assets do.

No single check steers a buildout, and none has to. It is enough that the one actor whose clock matches the system's clock invests in its own interest. The incentives align as rarely before: they align on the investor layer, and the window is open only until the pipeline calcifies around the alternative.

References

- ¹ Shehabi, Arman, et al. *2024 United States Data Center Energy Usage Report*. Lawrence Berkeley National Laboratory, 19 Dec. 2024, eta.lbl.gov/publications/2024-lbnl-data-center-energy-usage-report
- ² PJM Interconnection, 2027/2028 Base Residual Auction Report, Dec. 17, 2025, pjm.com
- ³ Environmental Integrity Project. *The Power Behind AI: Wave of Dirty Gas Power Plants Planned for Data Centers*. Environmental Integrity Project, 1 July 2026, environmentalintegrity.org/wp-content/uploads/2026/06/EIP_Report_ThePowerBehindAI_7.1.26.pdf
- ⁴ Union of Concerned Scientists. *Data Center Power Play: How Clean Energy Can Meet Rising Electricity Demand While Delivering Climate and Health Benefits*. Union of Concerned Scientists, Jan. 2026. *Union of Concerned Scientists*, www.ucs.org/sites/default/files/2026-01/Data-Center-Power-Play-report.pdf
- ⁵ Royal, H. (2026). The Compute Heat Rate: Quantifying AI-Driven Electricity Price Tolerance and Its Implications for Wholesale Market Repricing. SSRN 6322318. DOI: 10.2139/ssrn.6322318
- ⁶ Frontier Inference is flagship model API serving at premium pricing. Mid-Tier Inference covers mainstream production model serving. Enterprise Agentic AI is multi-step autonomous task execution for business workflows. Enterprise Contracted refers to committed-volume enterprise deployments at negotiated rates. Commodity Inference is generally open-weight and heavily discounted bulk serving
- ⁷ The blended CHR is the consumption-weighted average of the individual workload-tier CHRs, not the CHR formula applied to the blended R(w) and C values; applying the formula to the blended-row inputs yields ≈\$6,269/MWh, with the ~\$80/MWh difference reflecting the tier weighting
- ⁸ Siap, David. *Affordability Is All You Need: Meeting New Peak Capacity and the Next Era of U.S. Grid Investment*. CREO, Feb. 2026, http://www.creosyndicate.org/wp-content/uploads/2026/01/2026_CREO_Energy_Affordability.pdf
- ⁹ **Data sources.** Capacity additions are from utility IRPs and PJM planning data; capital and operating costs from the NREL Annual Technology Baseline and Lazard, with confidential utility cost decks proxied by these public estimates; ELCC values from PJM's published class ratings; and the cost-of-capital, tax, and rate-base parameters from utility regulatory filings
- ¹⁰ PJM Interconnection (2026). Powering Reliability Through Market Design. May 6, 2026
- ¹¹ Blackwell, Stan. "Data Center Forecasting." Presentation at the 2025 IRP Contemporary Issues Technical Conference, Indiana Utility Regulatory Commission, 18 Sept. 2025
- ¹² Mamkhezri, J., Sun, X., and Yang, C. (2025). The Hidden Cost of the Cloud: How Data Centers Affect Wholesale Electricity Prices. USAEE Working Paper No. 25-660, SSRN 5736562
- ¹³ The Fuel-Side Corollary extends the CHR framework from electricity into fuel markets: flat 24/7 data center load raises baseload gas demand, lifting gas prices, which feed back into the operating costs of gas-fired generation. The same demand growth that establishes the electricity-side price tolerance therefore erodes the gas ratio from below. A full treatment is in preparation
- ¹⁴ The self-reinforcing cycle in which rising wholesale prices push price-sensitive industrial load off the system over time, concentrating remaining demand among price-inelastic data centers, which supports further price increases
- ¹⁵ Voltus, Inc. "Voltus and Google to Deliver Grid Capacity and Local Economic Benefits Through Bring Your Own Capacity Agreement." Voltus, 2 June 2026, www.voltus.co/press/voltus-google-bring-your-own-capacity-pjm.
- ¹⁶ Rewiring America. "New Report: Household Upgrades Could Offset All New Projected Data Center Energy Demand." Rewiring America, 18 Sept. 2025, www.rewiringamerica.org/newsroom/press-releases/peak-demand-report-2025

-
- ¹⁷ CREO Syndicate. *Virtual Power Plants: Scaling Considerations in a Fragmented Landscape*. CREO Syndicate, May 2024
- ¹⁸ Noting that MCC is fundamentally a FTM metric given it has utility cost of capital as an input
- ¹⁹ The Environmental Integrity Project finds most of the 143 GW of planned BTM gas defaults to simple-cycle turbines rather than the more fuel-efficient combined-cycle class
- ²⁰ "NAACP Sues xAI for Illegal Pollution from Data Center Power Plant." *NAACP*, 14 Apr. 2026, naacp.org/articles/naacp-sues-xai-illegal-pollution-data-center-power-plant