ROYAL POWERS ENERGY

The grid has room. Does the market?

A practical framework for powering AI growth while protecting electricity customers



Cover image adapted from a photograph by American Public Power Association / Unsplash. Illustrative grid infrastructure.

John Powers
Royal Powers Energy, Inc.
November 3, 2026

Executive summary

AI data centers want to connect to power as quickly as possible. The economics help explain why: the Q3 2026 Compute Heat Rate estimates that an average new AI data center running inference could theoretically pay $5,631/MWh for electricity and still cover costs and margin. Getting online sooner can be extraordinarily valuable. But communities and policymakers are asking what that growth means for grid reliability, electricity bills and the people who host it. [3,17,25]

Flexible loads offer a promising way forward. Electric grids are built for their most difficult hours, leaving significant room during much of the year. A large load, like a data center, that reduces demand when the grid is under stress may connect sooner while new generation and transmission are built. Research from Duke and Brattle explains the opportunity. This paper asks what happens to electricity prices when that new demand runs during the rest of the year and no new supply is added to meet that demand in the short term. Physical headroom is not necessarily economic headroom. [1,2]

RP Energy tested hypothetical additions of new, around-the-clock load from 1 through 10 GW against ERCOT five-minute records from December 5, 2025 through October 4, 2026. We looked at the price impacts of adding the load with no flexibility, a “peak-neutral” case where the theoretical load curtails at the previously observed system peak (around 91.3 GW) and a “price-responsive” case where the new load also curtails when the modeled price reaches certain thresholds ($100/MWh, $150/MWh, $250/MWh).

At 1 GW of new load, the model estimates just under $1 billion in additional wholesale electricity costs for existing customers over the ten months, and as the load stack increases the curve steepens substantially. Adding 10 GW of new peak-flexible load requires curtailment of just 2.68% of that load to avoid a new system peak, but the model estimates roughly $118 billion in additional wholesale electricity costs for existing customers in ERCOT. Shifting to price-triggered curtailment reduces that wholesale cost increase at the same level of electricity consumption in all cases, with the largest differences at higher loads. The high-load assumptions are discussed in Section 3. [4-6]

While the study shows the significant benefit of load flexibility on market pricing, there remains a substantial increase in wholesale electricity costs for existing customers even after flexibility is applied. As most of the major data center operators and hyperscalers intend to protect other ratepayers, other solutions in addition to flexibility need to be considered. The most practical response combines three proposed actions: (1) make loads flexible, (2) bring new generation that reduces the cost of supplying the market, and (3) protect existing customers from remaining costs while allowing host communities to share in the economic opportunity. A well-designed system that combines all three approaches stands to allow for faster and more efficient connection of large loads while going beyond just “protecting” other ratepayers and instead directly benefiting them along with host communities.

Physical headroom is not necessarily economic headroom.

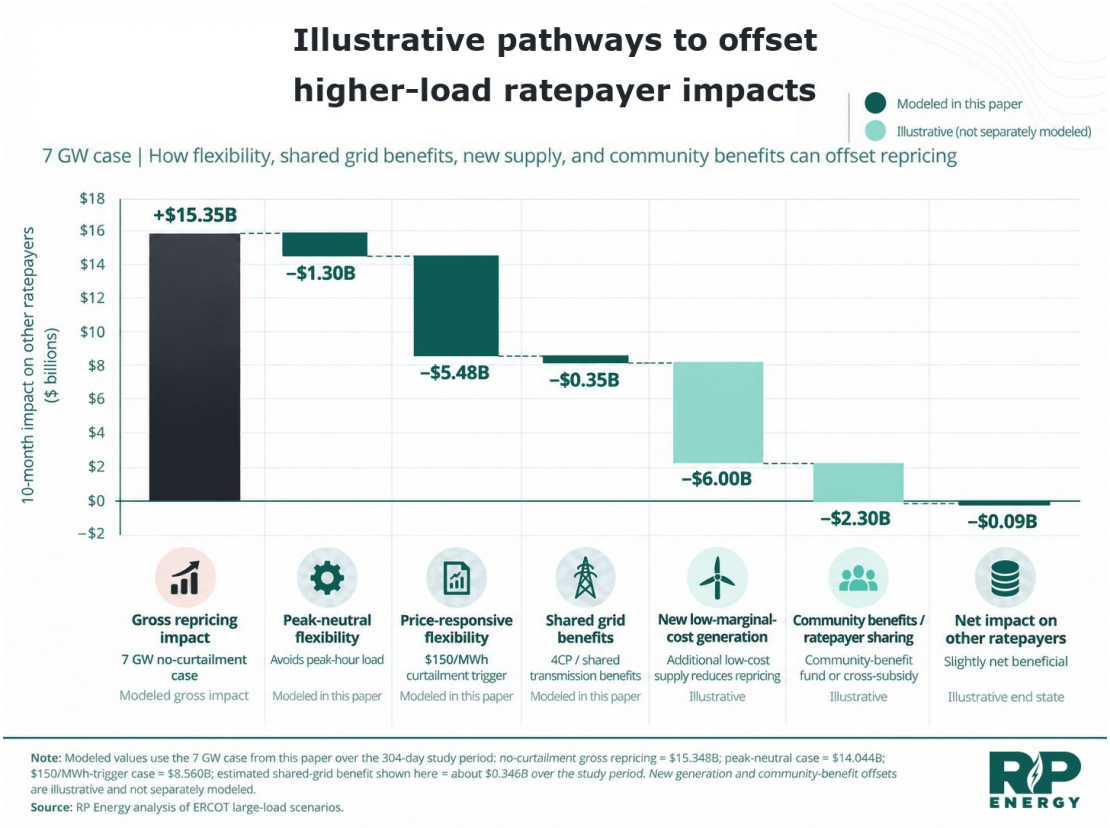
5 GW of new AI demand can operate below ERCOT's historical peak with less than 1% curtailment, and still add about $4.4 billion to wholesale electricity costs for existing customers over ten months. [4-6]

Table 1. Modeled wholesale electricity cost increases for existing customers grow rapidly as new load increases, while price-responsive flexibility becomes increasingly valuable

Added load	No curtailment	Peak-neutral	$150/MWh curtailment
1 GW	$0.891B	$0.891B	$0.877B
3 GW	$2.485B	$2.478B	$2.425B
5 GW	$4.439B	$4.395B	$4.207B
7 GW*	$15.348B	$14.044B	$8.560B
10 GW*	$144.253B	$118.209B	$18.826B

* Modeled wholesale electricity cost increases over ten months, in $ billions. The full 1-10 GW results are in Appendix Table A1. At 7 and 10 GW, some scenarios extend beyond the published supply-curve quantity. The treatment is explained below and in Appendix A6. [4-6]

Figure 1. Illustrative pathways to offset higher-load ratepayer impacts (7 GW case).



Scope of this study

We examine ERCOT on a system-wide level rather than modeling the load addition at an individual load zone, hub or node. In practice the location of the new load would have a significant impact on the results. Local transmission constraints and available generation could increase or reduce the price impact. Studying where on the grid new load can connect with the least effect on other customers is recommended for follow-on analysis.

1. Why speed to power matters

AI companies have strong incentives to build and energize data centers quickly. A new facility involves major capital commitments for chips, cooling equipment, networking and buildings. Delays in securing power can leave those assets underused while demand for AI services continues to grow. The economics are unusual because a megawatt-hour of electricity can support computing services worth far more than the price of conventional electricity consumption.

The Compute Heat Rate, published by RP Energy, calculates that as of Q3 2026 an average new AI data center running inference loads could tolerate electricity prices of $5,631/MWh and still cover costs and margin. Comparing this to wholesale electricity prices of roughly $30-100/MWh shows just how much headroom there is in the economics of turning electricity into compute.

Once you take fixed costs out of the decision and ask what an existing data center could pay for power before it stops making sense to run, that number jumps to $9,071/MWh. This is the Dispatch Compute Heat Rate, or CHR-D.

Using CHR-D, if an AI data center pays $70/MWh for power and runs at an 80% capacity factor, the illustration implies theoretical value creation of $63 million per installed MW per year. So for 1 GW of new data center load, being able to connect one year earlier could be worth on the order of $63 billion. [3,17]

Utilities and grid operators have good reasons to explore faster connection as well. Duke University research led by Tyler Norris estimated significant room for new demand that can turn down at critical times, including roughly 10 GW in ERCOT at a 0.5% average annual curtailment rate. Brattle examined how large loads could connect sooner by accepting limits on when they draw power. Both approaches make better use of existing infrastructure. Our study builds on that work by asking what the new load does to energy prices while it is running. [1,2]

Other research reinforces the need to combine flexibility with new supply. A Camus, encoord and Princeton ZERO Lab study in PJM found that both flexible service and developer-backed generation could reduce the cost of serving new data centers. The University of Texas large-load symposium describes a similar combination of flexible operation, timely infrastructure and new generation. [20,21]

Public expectations are changing at the same time. In March 2026, major technology companies accepted the White House Ratepayer Protection Pledge, including commitments to bring or buy new power, pay for grid upgrades associated with their projects and work with utilities on new rate structures. Communities are asking similar questions about local jobs, air quality, water use and their share of the economic benefits. Developers that can offer both a

credible power plan and credible customer protections have a better path to durable support. [25]

As the conversation moves from reliability to affordability, the details of those agreements become critical. A developer may pay for its own substation, contribute to existing wires and promise to turn down during emergencies. The remaining question is what happens to the price of electricity during the many hours when the facility runs normally.

2. How adding load can change electricity prices

Electricity markets generally call on lower-cost generation before higher-cost generation. A wind or solar project may offer electricity at a very low price because it has no fuel cost. A gas plant must cover fuel and other operating costs. As demand grows, the system works its way up this supply stack. The price needed to attract the next generator can then set the price for a much larger volume of electricity already being sold. [14]

Throughout the paper, when we refer to higher wholesale electricity costs for existing customers, we mean the additional wholesale spending on electricity ERCOT customers were already consuming as a result of the modeled price increase. How much of that ultimately reaches individual customer bills depends on contracts, hedges and retail rate structures. Appendix A refers to this technical measure as modeled wholesale repricing.

A real ERCOT interval illustrates the mechanism. At 02:55 on July 13, 2026, the historical marginal offer was $24/MWh. Adding 5 GW moved the modeled marginal offer to $28/MWh, a $4/MWh increase. Applied across 55.4 GW of existing ERCOT load in that five-minute interval, that produces about $18,462 in additional wholesale electricity costs. Across all 288 five-minute intervals that day, the 5 GW peak-neutral case adds about $5.66 million. Figure 2 shows the same offer curve steepening further at higher theoretical load additions. [4-6]

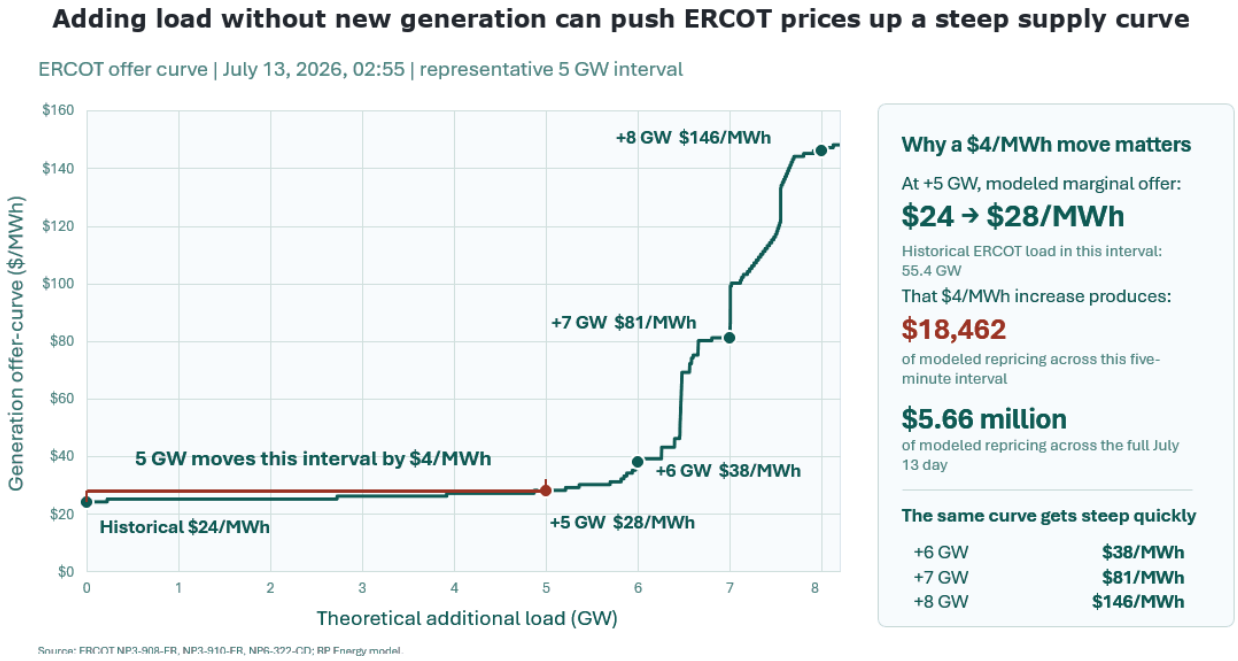


Figure 2. July 13, 2026 at 02:55: adding load moves the modeled marginal offer up ERCOT's disclosed generation offer curve. In this representative interval, 5 GW of added load increases the modeled marginal offer from $24/MWh to $28/MWh; larger additions move farther up the same curve. [4-6]

There is an important benefit on the other side of the ledger. New customers help pay for power lines, substations and other infrastructure that has already been built. Spreading those fixed costs across more electricity sales can reduce some cost burdens for existing customers. Berkeley Lab and Brattle have shown that load growth can be associated with either higher or lower retail rates depending on infrastructure utilization, new supply, market conditions and tariff design. [19]

Our transmission-sharing estimate shows how much this could help. At 1 GW, the new peak-neutral load would contribute roughly $63 million a year toward existing transmission costs. Adjusted to the ten-month study period, that is about $53 million, compared with $891 million in added wholesale electricity costs for existing customers. At 3 GW, the corresponding amounts are about $149 million and $2.48 billion. The contribution matters, but amounts to only about 6% of the wholesale cost increase at these load levels. Distribution charges require a site-specific estimate. [28]

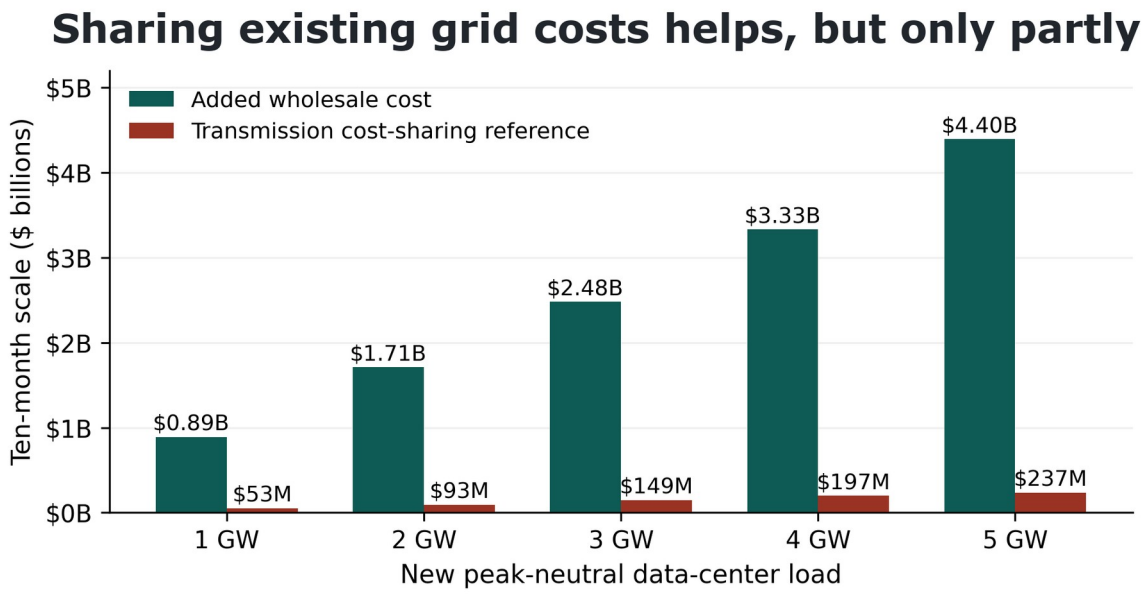


Figure 3. Peak-neutral increase in wholesale electricity costs for existing customers compared with the new load's estimated contribution toward existing transmission costs. Both are scaled to the 304-day study. Distribution benefits are not estimated. [28]

When added load raises wholesale market prices, some of the additional spending is transferred from buyers to sellers. Generators selling electricity at market prices may earn higher revenues, while retail suppliers, industrial customers and households can be exposed to higher costs depending on contracts, hedges and pass-through. Owners of virtual power purchase agreements can also receive higher payments when wholesale prices rise. The policy challenge is to identify who is exposed and how to protect them. [4-6]

3. What we tested in ERCOT

We modeled demand additions in every 1 GW increment from 1 through 10 GW, without adding new generation, storage or transmission in response. While in reality there will of course be generation added, there is often a lag in timing between the speed at which a new large load

can come online and the speed at which the new generation can be built and connected to serve it. This study models what a 10 month generation lag to load would have looked like since last December. The study covers December 5, 2025 through October 4, 2026. December 5 marks the start of ERCOT's Real-Time Co-optimization plus Batteries system, so all the data come from the same operating regime. The analysis uses 86,076 five-minute intervals with matched load and pricing data. Appendix A explains the data and calculations. [4-6,24]

Three operating cases are compared at each load level. Unrestricted load draws its full electricity demand in every interval. Peak-neutral load backs down only enough to avoid exceeding ERCOT's highest observed system demand, about 91.3 GW. The price-responsive case (run at price thresholds of $100, $150 and $250/MWh) first applies the peak-neutral cap and then reduces load further when necessary to hold the model's marginal offer price at or below the selected threshold.

In the 5 GW case, peak-neutral operation uses roughly 99.1% of the electricity the new load would consume if it ran continuously (less than 1% curtailment), yet adds $4.395 billion to wholesale electricity costs for existing customers. The $150/MWh price rule lowers that to $4.207 billion while serving 98.8% of potential demand. That is a reduction of roughly $188 million, with more than $4 billion still remaining. At 1, 2 and 3 GW, the peak-neutral wholesale cost increases are $0.891 billion, $1.711 billion and $2.478 billion, respectively. Appendix Table A1 includes every 1 GW increment. [4-6]

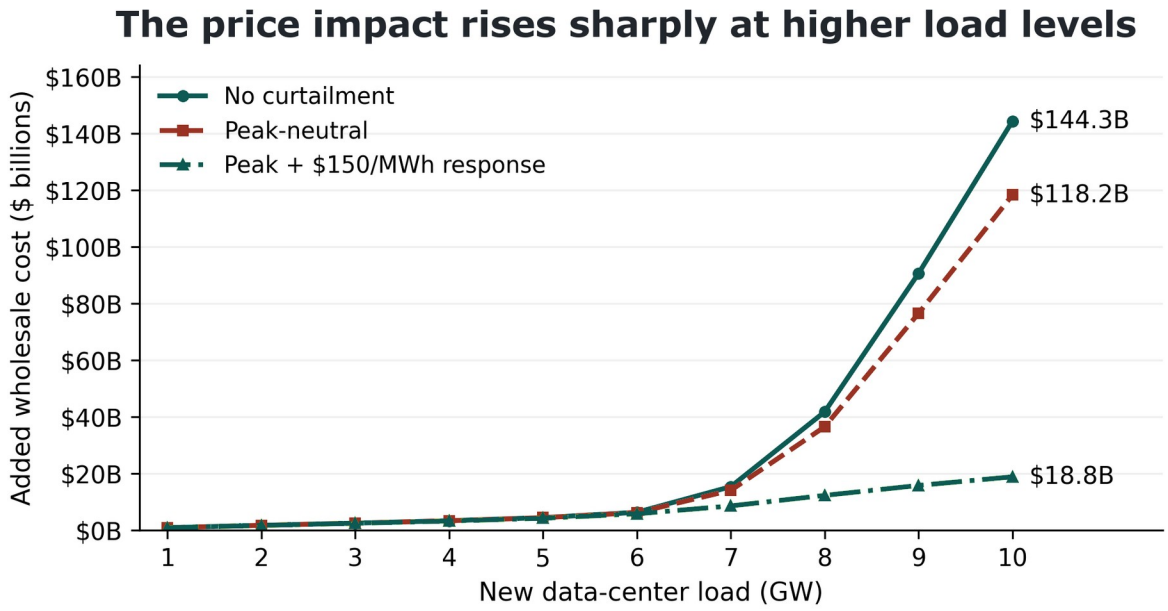


Figure 4. All three operating strategies at every 1 GW increment. The wholesale cost impact rises sharply above 6 GW. [4-6]

From 7 to 10 GW of additional load, the modeled increase in wholesale electricity costs for existing customers jumps significantly, and the impact of curtailing at a certain price threshold is most pronounced. Adding 10 GW without curtailment produces over $144 billion in additional wholesale electricity costs in the model. Curtailing at historical system peak drops this to $118 billion, and curtailing above $150/MWh reduces it to $18.8 billion (with roughly 6% curtailment).

When adding significant theoretical loads to the historical supply stack, there are times when the new load exceeds the amount offered by generators in a particular five-minute interval. In a real market, more generators might have offered power if that additional demand were expected. For this exercise, we held the last offered price constant for load above the observed generation offer stack. The larger-load results therefore depend more on this assumption, which Appendix A6 explains in detail. [4-6]

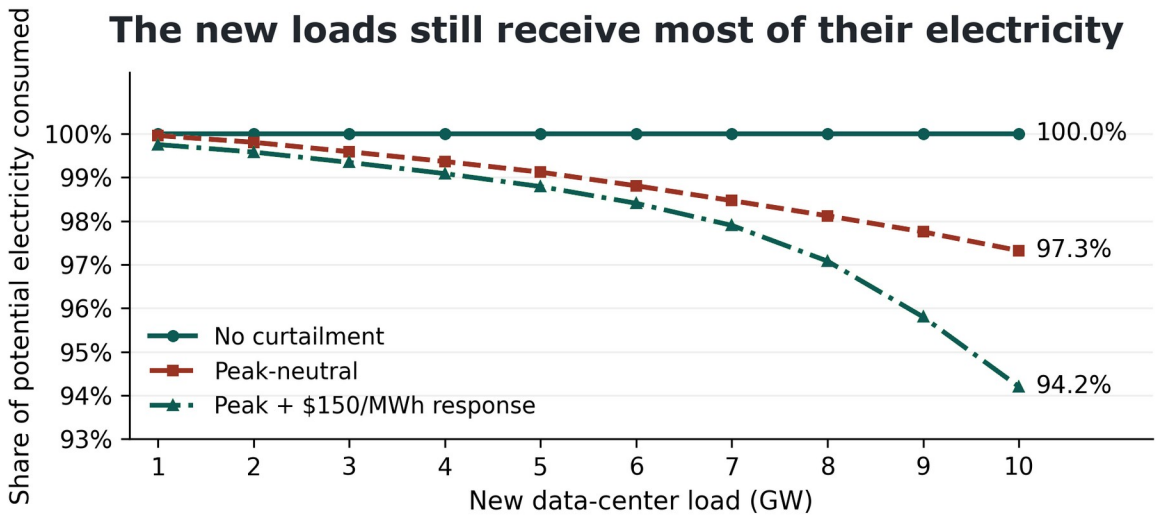


Figure 5. Share of potential electricity demand served under all three operating strategies. Even at 10 GW, the $150/MWh case serves about 94% of potential demand. [4-6]

To compare the different curtailment strategies, we also looked at the wholesale cost impact at a common curtailment percentage. In the 5 GW case, when we hold the curtailment percentage constant, the $150/MWh price rule produces about $175 million less added wholesale cost than peak-based curtailment, a 4% improvement. Appendix A5.2 contains the chart and all 30 comparisons. [4-6]

Based on this analysis, adding significant new load without new generation is likely to increase wholesale electricity costs for existing customers. What are the steps new large loads could take to mitigate and even reverse this effect?

4. Step one: Offer dependable flexibility

The first step is to make flexibility part of the connection agreement. A data center should be able to reduce its draw from the grid when the system needs relief. It can also turn down during expensive hours, even when total demand is well below the annual peak. Our results show why both matter: peak-based flexibility makes room on the system, while price-based flexibility can reduce the cost of serving the new load.

Flexibility can be provided in several ways. A data center can temporarily curtail computing workloads, shift work to another location or time, draw from a charged battery, or use permitted generation behind its meter. A virtual power plant (VPP) can coordinate batteries and other customers to provide relief elsewhere on the grid. The practical question is which option delivers the required reduction in grid demand, for the required duration, at the lowest cost. Location matters when that relief is provided somewhere else.

CHR-D helps explain why operators may prefer to keep the computers running. It is the break-even electricity price for running an AI data center once fixed costs are taken out of the decision. At what price does it no longer make sense to turn electricity into AI compute? The Q3 2026 figure is about $9,071/MWh. That creates a strong incentive to use batteries, VPPs or on-site power rather than stop valuable computation, where those alternatives can deliver the required flexibility. [3,17]

This is also where policy design can create unintended consequences. If developers are required to stop using grid electricity during tight periods, they may have a strong incentive to keep the computers running on backup diesel generators (assuming batteries, VPPs or other solutions are not available). Emergency engines face permitting and operating restrictions, and extensive use could create local air quality problems. Batteries and other cleaner forms of flexibility should be encouraged where they can meet the required duration and reliability standards. The objective is to reduce strain on the grid without creating a new problem for the community. [26]

5. Step two: Bring generation that lowers the cost of power

At 5 GW, requiring the new load to curtail at $150/MWh still adds about $4.2 billion to wholesale electricity costs for existing customers over ten months. Flexibility therefore needs to work alongside new supply. Adding a lower-cost generator (whether physical or in the form of a virtual power plant) can move the market in the opposite direction, reducing the price of electricity for the new load and other customers.

Developers are increasingly committing to bring power with their projects. The most useful question is whether that supply actually changes the cost of electricity when and where the system needs it. A generator that produces exactly the same number of megawatt-hours as a data center consumes may still leave expensive hours uncovered. Conversely, inexpensive generation that is available during hours when the market would otherwise call on expensive generators could have an outsized benefit even if its hourly output is not a perfect match for the facility. Location, transmission capacity and the timing of delivery all matter. The same framework applies whether generation is built behind the meter or connected elsewhere on the grid.

There are already constructive examples from many data center developers and hyperscalers. Looking just at Google, we see several such capacity plays in the recent news. In October 2026, Google and Constellation announced an agreement intended to add 890 MW of nuclear output to the PJM system by upgrading existing nuclear units. In Michigan, Google announced plans with DTE to support 2.7 GW of new clean resources, including solar, storage and flexible demand, and a $10 million energy impact fund. Michigan regulators subsequently approved special contracts with provisions intended to protect other customers from the project's costs. On October 22, 2026, Google and Voltus announced an expanded agreement for a roughly 1 GW virtual power plant in PJM. Google will fund a portfolio of batteries, smart thermostats and flexible industrial equipment at homes and businesses, with participating customers receiving bill savings or direct payments. The agreement is a practical example of a data-center company paying for flexibility elsewhere on the grid while sharing the benefits with other electricity customers. [31,32] These are examples of developers combining new supply, customer protection and community benefits. [8-10,27]

Texas is already debating how to combine flexible service, staged connections and developer-backed power supply. Discussions around ERCOT's Batch Zero process and the University of Texas large-load symposium reflect the urgency of getting those rules right. Our analysis adds a reason to include affordability alongside reliability when designing them. [18,21]

Low operating cost and low emissions are especially valuable goals. Solar, wind and nuclear power can supply energy without burning gas during every hour of operation, although each technology has different availability and investment characteristics. Storage can move energy from lower-cost hours into more expensive hours. Dispatchable gas generation may still have a place in a broader portfolio, but a large wave of gas-fired projects would also increase demand for fuel and the equipment needed to build and connect those plants. Those effects deserve attention alongside any on-site reliability benefit.

6. Step three: Protect ratepayers and share the economic upside

Flexibility and new supply should reduce the pressure on electricity prices, but some cost increases may remain. A data center developer could cover that cost through a dedicated tariff, bill credit, utility agreement or another transparent payment. Requiring the new load to take responsibility would give it a reason to find the least expensive combination of curtailment, storage, new generation and purchased flexibility. The agreement would need to identify which costs actually reach existing customers.

There are two useful ways to express the wholesale cost increase in dollars per MWh. The first asks how much the average wholesale cost rises across the electricity existing ERCOT customers use. At 5 GW, peak-neutral operation adds about $10.35/MWh. The $150/MWh rule lowers that to about $9.90/MWh. How those wholesale changes reach individual bills depends on customers' contracts and rates. [4-6]

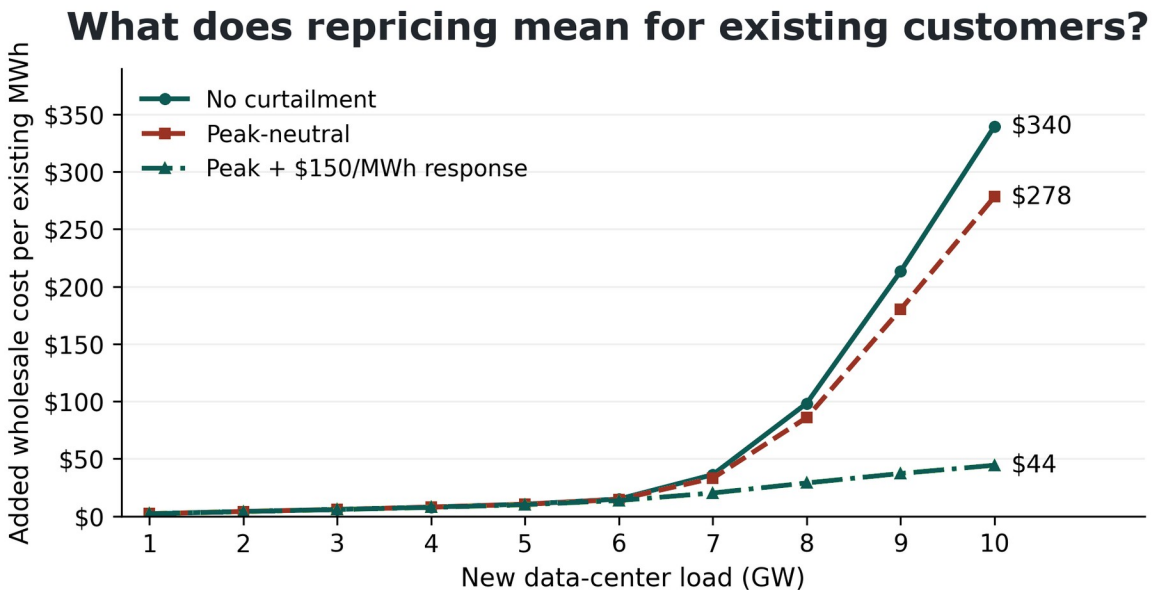


Figure 6. Modeled increase in wholesale electricity costs per MWh of existing ERCOT consumption. All three cases use the same 424.77 million MWh denominator. [4-6]

The second asks what the new load would pay per MWh of its own consumption to cover the modeled increase in wholesale electricity costs. At 5 GW, the amount is about $124/MWh for peak-neutral operation and $119/MWh when the load curtails at $150/MWh, in addition to the cost of buying its electricity. At 10 GW, the corresponding amounts rise to about $1,694/MWh and $279/MWh. This illustrates the size of a full-offset payment. An actual arrangement would account for who bears the cost and what other benefits the project provides. [4-6]

The $119-$124/MWh payment in our 5 GW cases is roughly 2% of the $5,631/MWh Q3 CHR estimate discussed earlier. For a developer that can earn far more by getting its computers running sooner, covering that extra cost could make commercial sense. It could also bridge the gap between a data center opening and the new generation intended to serve it coming online, with the payment falling as that supply reduces the impact on existing customers. [3,4-6,17]

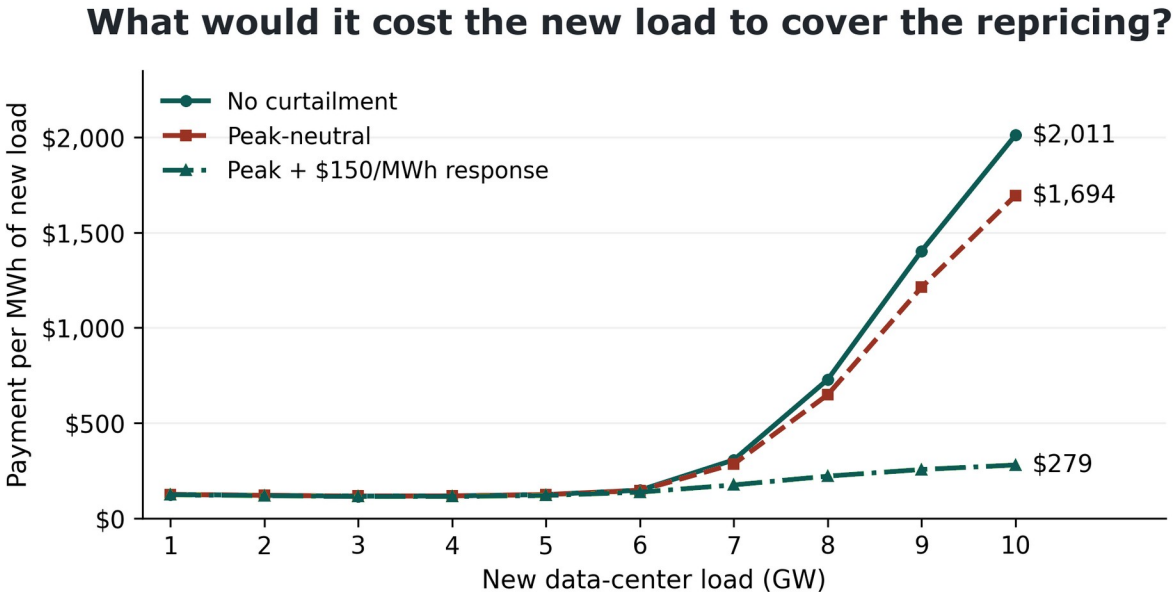


Figure 7. Payment per MWh consumed by the new load to offset the full modeled increase in wholesale electricity costs. These are additional to electricity procurement costs. [4-6]

Community benefits are a further opportunity. A developer can protect utility customers from additional energy costs and also help the people living near the site share in the value of the investment. That could include a fund for household efficiency improvements, direct local support, workforce programs or a recurring community dividend. One model worth exploring resembles the Alaska Permanent Fund, which allows residents to participate in economic value generated by natural resources. Recurring contributions could be invested, with a share of the fund paid to residents over time. The goal would be a durable arrangement that residents understand and trust, rather than a one-time announcement at the beginning of a project. [11]

One future application for CHR would be to link those contributions to a small agreed fraction of the published National CHR multiplied by metered data-center electricity use, settled quarterly. The payment would then move with the economics of AI rather than depend on a company disclosing profit at each facility. PJM's May 2026 market-design paper cited CHR in its discussion of large-load demand flexibility and scarcity pricing, although it did not propose a

community fund. A host community would determine how any fund is governed and who benefits. [3,11,17,30]

Even a small fraction could be meaningful. Using the Q3 2026 CHR reference, an illustrative contribution of 0.5% is about $28/MWh. For a 1 GW facility running nearly continuously, that is roughly $247 million a year. To make the comparison useful, Table 2 expresses annual support per household using an illustrative 100,000-household community for the total-dollar examples. [17,18]

Table 2. Annual community-benefit scale, using 100,000 households for the total-dollar examples. [17,18]

Community-benefit reference	Annual equivalent per household
Indiana household rate proposal	$100 per year (reported proposal)
Alabama: $270 million over 20 years, averaging $13.5 million a year	$135 per year
0.5% of Q3 CHR; hypothetical 1 GW load contributing $247 million a year	$2,470 per year

The common household count illustrates scale, not the actual populations or payouts of these projects. Pennsylvania's reported $10,000-per-household proposal is a one-time payment, so it is kept outside the annual comparison. A permanent fund could invest contributions rather than distribute the full amount. [18]

Ratepayer protection addresses the electricity costs a project creates. A community fund lets residents share in the broader value of hosting the industry. Combining the two could give developers a clearer path to approval and give communities a reason to welcome the investment.

7. What else changes as we add AI load?

The power-price effect is only one consequence of rapidly growing electricity demand. New demand can also change which generating plants run and therefore the amount of carbon emitted. If extra electricity comes from higher-cost fossil units, that can raise both fuel consumption and emissions. New low-emission supply, virtual power plants and batteries can reduce those effects, depending on how they operate. Quantifying those emissions, including local air-quality impacts from backup engines, is a recommended scope for a follow-on study. [12,13]

The wider energy system matters as well. New gas-fired generation creates additional demand for natural gas, which can influence fuel costs for other generators and consumers. Transmission lines, transformers, turbines, conductors and switchgear compete for manufacturing capacity and skilled labor. The U.S. Department of Energy has highlighted long lead times for critical grid equipment and the pressure that growing electricity demand puts on those supply chains. The more effectively we can use existing infrastructure and bring genuinely incremental supply online, the better we can manage these costs. [22,23]

Using existing grid capacity more efficiently could also delay or avoid future upgrades. That benefit is separate from sharing the cost of infrastructure already built. Appendix Figure A2 illustrates the possible scale, but a local planning study would need to identify the actual projects that could be deferred. [28,29]

The same questions arise outside Texas, even though the markets work differently. In PJM, large loads can affect both energy prices and payments for generation capacity. In a regulated utility territory, the cost may reach customers through fuel charges, purchased power and later

rate decisions. The common question is how to connect new demand quickly while protecting the customers already on the system. [15,16,19]

8. A better bargain for AI and electricity customers

AI data centers will continue seeking electricity because the economic opportunity is substantial. The question is how to turn that investment into a stronger grid and a better deal for existing customers. Our ERCOT analysis shows that 5 GW of new load can avoid setting a new system peak with less than 1% curtailment, while still adding about $4.4 billion to wholesale electricity costs for existing customers over ten months. Price-responsive flexibility improves the result, but does not remove the need for other solutions.

A well-designed approach combines dependable flexibility, new low-cost generation, protection from remaining electricity costs and meaningful community benefits. It gives developers more ways to connect sooner while helping ratepayers and host communities share in the opportunity. The next step is to apply that approach to specific locations and agreements, where the right combination can be tested and delivered.

We welcome feedback and collaboration

This paper is meant to be a practical starting point. We would welcome feedback on the assumptions, methods and implications, and we would be glad to discuss how the framework could be applied to particular utilities, load zones, projects, tariffs or policy questions.

If you would like to compare notes, explore a joint study or discuss a specific application, please visit rpenergyinc.com or contact John Powers at john@rpenergyinc.com.

Appendix A. Methodology, sources, and interpretation

A1. Purpose and study window

This historical counterfactual examines the short-run wholesale price response to incremental ERCOT data-center demand without adding generation. It covers December 5, 2025 through October 4, 2026, inclusive, at five-minute resolution. December 5 is the effective starting date because ERCOT put Real-Time Co-optimization plus Batteries (RTC+B) into operation that day; earlier offer data come from a different real-time market regime. The model tests 1 through 10 GW under unrestricted, peak-neutral and $150/MWh price-responsive operation. The follow-up compares matched-energy physical caps with $100, $150 and $250/MWh price rules at each GW level. The method is a static historical stack counterfactual, not a long-run dispatch, security-constrained nodal or retail-rate forecast. [4-6,24]

A2. Input datasets and observation coverage

Historical load and dispatch reference information come from ERCOT NP3-910-ER. Aggregate non-intermittent-resource offers come from NP3-908-ER; SCED System Lambda provides historical market-price context. Across 87,530 usable five-minute load observations, 86,076 intervals have the required matched priced inputs. Five days of offer-curve observations and 14 other intervals without System Lambda are excluded. The complete 1–10 GW analyses and the matched-CF comparisons use the same 86,076 priced intervals. [4-6]

The full load time series is used for dispatch energy and historical-peak exceedance calculations. Wholesale repricing is calculated only for intervals with sufficient joined pricing data. A total of 1,454 load intervals lack joined prices. These comprise five excluded days without usable offer-curve coverage (March 4, May 3, June 18, August 1 and September 21, 2026) and 14 intervals without matching System Lambda observations. Missing observations are not filled with synthetic prices. For price-responsive capacity-factor calculations, the model retains peak-neutral dispatch without additional price-triggered curtailment during unpriced load intervals.

A3. Aligning published offer curves with historical dispatch

ERCOT's NP3-908 offer curve contains the disclosed economic offers for non-intermittent resources, while NP3-910 also reports resources operating on Output Schedules. To represent both on a common quantity axis, the model shifts each NP3-908 offer-curve MW point by the relevant Output Schedule MW. The historical dispatch reference retains Base Point non-IRR plus Output Schedule. This is equivalent, for calculating an incremental price change, to leaving the economic curve unshifted and locating the historical reference on the Base Point non-IRR axis. [4,5]

For any modeled demand quantity, the lookup uses the first disclosed offer price whose cumulative quantity is greater than or equal to the target. The baseline is the offer-curve price at historical dispatch and the counterfactual is the price at dispatch plus incremental served load. The nonnegative difference provides the repricing delta and System Lambda serves as a historical market-context anchor. The October 24, 2026 raw-curve replay examined 360 reports, 254.4 million offer rows and 86,076 selected priced intervals, reproducing the originally reported scenario gross values. [4-6]

A4. Three operating strategies

In the no-curtailment case, each modeled load level adds its full nameplate demand in every five-minute interval, producing a 100% energy-based capacity factor by assumption. The model separately calculates the number of historical intervals in which total system demand would exceed the observed 91.308 GW peak, the largest exceedance and the total demand energy above that peak.

In the peak-neutral case, served additional MW in each interval is the smaller of nameplate demand and the remaining headroom between historical system load and the observed 91.308 GW peak. If the historical system load is at or above that limit, the new load is completely curtailed in that interval. Partial load is allowed when some headroom remains. This rule measures aggregate historical peak neutrality only.

The price-responsive case begins with the peak-neutral served quantity. It then reduces that quantity only as much as needed to keep the modeled offer-curve price at or below the chosen threshold of $100, $150 or $250/MWh in a priced interval. If the historical offer-curve anchor is already above the threshold, no incremental demand is served in that interval. A numerical search finds the maximum load that remains within the modeled threshold. The paper uses the $150/MWh case for the main comparison. These rules assume responsive, five-minute operation without minimum shutdown times, ramp constraints or storage energy limits.

A5. Gross modeled wholesale repricing

For each matched interval, the model multiplies the historical ERCOT load in megawatt-hours by the nonnegative modeled increase in the aggregate offer-stack marginal price, expressed in dollars per megawatt-hour. A five-minute MW value is divided by 12 to convert it to MWh. The sum of those products across 86,076 matched intervals is the gross wholesale repricing number reported in the main text. The historical ERCOT demand represented in those priced intervals is approximately 424.8 TWh. [4-6]

The existing-load denominator is 424,768,667 MWh. Gross wholesale repricing measures the modeled increase in wholesale expenditure on that demand. It includes transfers between market participants and is not a calculation of incremental production cost, net social cost or retail bill changes. [4-6]

In reality, hedging, contracting, retail tariffs, nodal congestion, unit commitment, storage, reserves, generation additions, and gas-price changes affect realized market clearing and financial incidence. Some wholesale repricing transfers value from electricity buyers to unhedged generators and other long wholesale positions, including certain virtual power purchase agreements. A make-whole policy would need an independently justified customer-cost measure, rather than mechanically charging the full modeled gross amount.

A5.1. Complete 1 to 10 GW scenario results

Appendix Table A1. Ten-month modeled gross wholesale repricing, in $ billions. [4-6]

Added load	No curtailment	Peak-neutral	$150/MWh rule
1 GW	$0.891	$0.891	$0.877
2 GW	$1.713	$1.711	$1.681
3 GW	$2.485	$2.478	$2.425
4 GW	$3.347	$3.330	$3.230
5 GW	$4.439	$4.395	$4.207
6 GW	$6.309	$6.177	$5.770
7 GW	$15.348	$14.044	$8.560
8 GW	$41.749	$36.471	$12.272
9 GW	$90.577	$76.482	$15.777
10 GW	$144.253	$118.209	$18.826

Appendix Table A2. Utilization, historical-peak exceedance and offer-stack limits. [4-6]

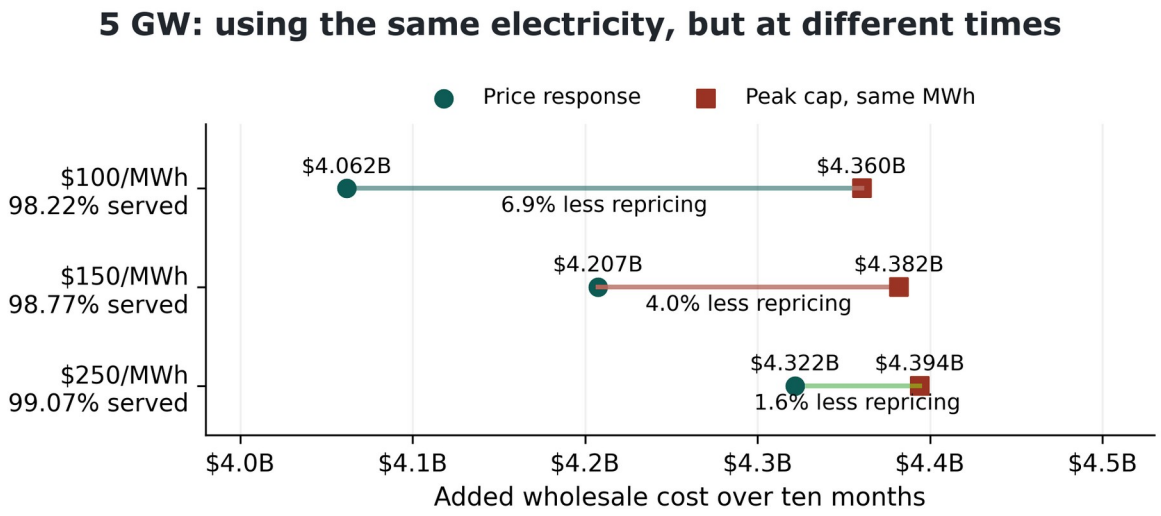
Added load	Peak-neutral CF	$150 rule CF	Peak excess hours, NC	Curve shortfall intervals, NC
1 GW	99.957%	99.751%	9.9	0
2 GW	99.803%	99.578%	44.9	0
3 GW	99.586%	99.342%	76.7	0
4 GW	99.363%	99.085%	114.5	0
5 GW	99.118%	98.788%	161.7	0
6 GW	98.806%	98.405%	234.5	0
7 GW	98.466%	97.898%	288.0	57
8 GW	98.118%	97.079%	341.4	856
9 GW	97.749%	95.800%	424.3	2,971
10 GW	97.320%	94.203%	532.9	6,923

Notes: NC is unrestricted operation. Historical-peak exceedance is a reference metric, not a forecast of outages. A curve shortfall means modeled demand exceeds the published offer quantity, not that ERCOT physically lacks generation. [4-6]

A5.2. Equal-electricity comparisons for price and physical flexibility

For every 1 GW increment, the follow-up calculates the incremental MWh that would be served if the price-responsive case applies $100, $150 or $250/MWh as its post-load offer-price threshold on top of the peak-neutral reference. A numerical solver then adjusts the physical system-load ceiling to serve precisely those same incremental MWh in the same 86,076 priced intervals. Gross repricing is evaluated for both dispatch profiles using the corrected raw offer curves. This isolates the timing of curtailment from the quantity curtailed. The target capacity factors below refer to priced intervals and can differ slightly from the whole-study capacity factors in Appendix Table A2. [4-6]

Price-responsive operation produces less repricing in all 30 matched-consumption comparisons. At 5 GW, the reductions are 6.9%, 4.0% and 1.6% at the $100, $150 and $250/MWh thresholds, respectively. The larger differences at 7-10 GW are subject to the offer-quantity limitations in Appendix A6. [4-6]



Appendix Figure A1. The 5 GW comparison at equal electricity consumption. Each pair consumes the same MWh but turns down at different times. The price-based rule produces less wholesale repricing at all three thresholds. [4-6]

Appendix Table A3. Percent reduction in gross repricing using a price-responsive rule rather than an equal-MWh physical load cap, and the priced-interval CF for the $150 rule. All 30 cases calculated. [4-6]

Load	$100 rule	$150 rule	$250 rule	$150-rule CF
1 GW	3.2%	1.4%	0.0%	99.747%
2 GW	3.8%	1.6%	0.1%	99.571%
3 GW	4.6%	2.0%	0.2%	99.331%
4 GW	5.6%	2.8%	0.7%	99.070%
5 GW	6.9%	4.0%	1.6%	98.768%
6 GW	8.6%	5.6%	3.3%	98.381%
7 GW*	40.5%	37.4%	31.8%	97.871%
8 GW*	67.4%	63.9%	57.6%	97.044%
9 GW*	80.0%	76.8%	70.5%	95.750%
10 GW*	84.3%	81.3%	75.3%	94.141%

A6. Observed offer-stack boundaries

For unrestricted or peak-neutral operation, hypothetical new demand can sometimes move beyond the maximum cumulative quantity disclosed in ERCOT's aggregate offer curve. The model holds the highest observed offer price when the target quantity exceeds that boundary and records the affected intervals and additional MWh separately. That price-holding rule is an explicit convenience for a static counterfactual, not a lower or upper bound on the price that would emerge with new commitments, generation, reserves and market response. The fact that an offer curve ends does not imply an actual generation shortage. [4-6]

The October 24 full-curve review quantified each strategy for all ten GW levels. All 1–6 GW scenarios remain inside the published offer-quantity boundaries. At 7 GW the peak-neutral scenario has 29 intervals past the boundary, with 2.8% of its gross repricing occurring in those intervals; at 8 GW the share is 23.7%, at 9 GW 37.0%, and at 10 GW 55.2%. At 10 GW, peak-neutral operation goes past disclosed quantities in 5,591 priced intervals and 0.643% of incremental MWh served lies beyond the quantity boundary. Its $150/MWh case has zero such intervals. Importantly, the 55.2% measures the total repricing occurring in affected intervals, not the dollars mathematically attributable to unsupported MWh. [4-6]

Appendix Table A4. Observed offer-quantity boundary for 1–10 GW. The last column counts the $150-rule intervals beyond the published curve. NC is no curtailment, PN is peak-neutral, only curtailing at the previously-observed system peak. [4-6]

GW	NC intervals	PN intervals	PN gross in these intervals	PN new MWh beyond curve	$150 intervals
1	0	0	0.0%	0.000%	0
2	0	0	0.0%	0.000%	0
3	0	0	0.0%	0.000%	0
4	0	0	0.0%	0.000%	0
5	0	0	0.0%	0.000%	0
6	0	0	0.0%	0.000%	0
7*	57	29	2.8%	0.001%	0
8*	856	711	23.7%	0.039%	0
9*	2,971	2,393	37.0%	0.217%	0
10*	6,923	5,591	55.2%	0.643%	0

A7. Historical peak and reliability

The observed historical peak serves as a reference point for the peak-neutral operating strategy. The 91.308 GW reference is the highest five-minute system load in the study data, recorded on July 22, 2026. ERCOT's published peak-demand record for that day is 91,089 MW, which ERCOT lists as unofficial until final settlement and reports on an hourly rather than five-minute basis. [5] In the ten-month 10 GW unrestricted case, total demand would exceed that historical peak for approximately 532.9 hours. In the peak-neutral and price-responsive cases, the model imposes no exceedance of the historical aggregate peak by construction. Actual grid reliability depends on sufficient generation and reserves, deliverable transmission service, local congestion, ramping, frequency response and security under contingencies. The model does not evaluate those impacts. Exceeding a historical peak is a reason to study reliability, rather than by itself proof of instability.

A8. Treatment of new supply, other costs and emissions

The modeled scenarios do not add wind, solar, nuclear, gas generation, batteries or virtual-power-plant supply in response to new demand. Nor do they simulate construction schedules, capital costs, fuel-price changes, unit commitment, reserve dynamics or marginal air emissions. A 4CP transmission cost-sharing sensitivity is calculated separately using modeled summer coincident load and a published cost-allocation reference; Appendix Figure A2 is a separate conditional infrastructure-deferral illustration. Neither estimate is netted from modeled wholesale repricing, and the former carbon screening calculation is not carried forward. [12,13,22,23,28-29]

A9. Geographic and market-design limitations

The ERCOT calculation is systemwide. It does not price a specific new data center at a load zone, node or transmission-constrained location. In practice, a low-cost generator near one project may not be able to serve another project behind congestion, and prices will differ across settlement points. ERCOT also makes dispatch and pricing decisions through a security-constrained market that is more complex than a simple aggregate stack. Price-triggered load operations would need an observable local reference, realistic notice and response timing, verification and settlement rules. Separate simulations are required for PJM capacity-market impacts or the fuel-cost and rate-treatment mechanisms of vertically integrated utilities. [14-16]

A9.1. Illustrative 4CP allocation of existing transmission costs

Transmission charges can reallocate an existing revenue requirement among ERCOT customers; this is not automatically an avoided cost of transmission construction. The October 24 follow-up applies the annual 2025 ERCOT 4CP reference of $68.55/kW-year to the hypothetical data-center net demand during modeled four summer 15-minute coincident peaks. It uses SCED load as a timing proxy, not official ERCOT settlement 4CP. At 5 GW peak-neutral operation the selected modeled periods imply an average 4,146 MW contribution, or approximately $284 million/year in referenced transmission allocation. Alternative tied summer peak selections yield about $174M to $337M/year, and the equivalent ten-month scaling is about 5.4% of that case's modeled gross wholesale repricing. Annual allocations are multiplied by 304/365 when compared with the 304-day study. These values are illustrative developer-side allocations, not realized credits to existing customers. Distribution sharing remains unestimated and may be zero for transmission-connected facilities. [28]

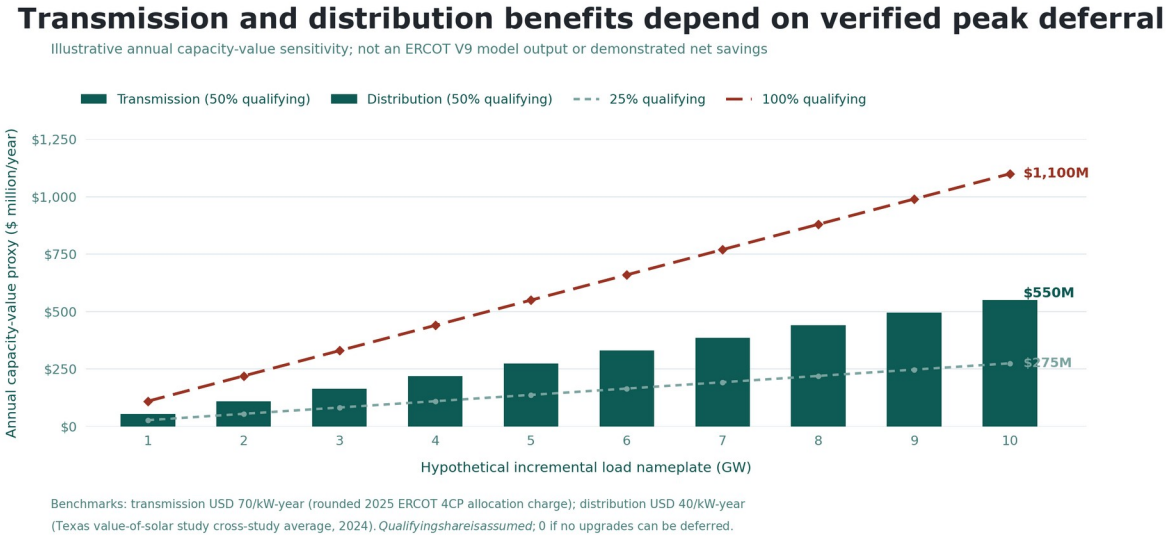
Appendix Table A5. Annual transmission cost of service allocation reference by load and operating strategy, in $ millions. The final column is the selected peak-neutral contribution period-adjusted to the ten-month gross repricing denominator. [28]

GW	Unrestricted	Peak-neutral	$150 price	PN allocation/gross
1	$69	$63	$63	5.9%
2	$137	$112	$112	5.4%
3	$206	$178	$178	6.0%
4	$274	$236	$236	5.9%
5	$343	$284	$284	5.4%
6	$411	$320	$320	4.3%
7	$480	$416	$416	2.5%
8	$548	$492	$492	1.1%
9	$617	$566	$566	0.6%
10	$686	$613	$613	0.4%

A9.2. Conditional capacity-deferral sensitivity

Appendix Figure A2 uses an independent annual capacity-value illustration rather than the wholesale model or the 4CP cost-sharing calculation. For G GW of added nameplate and qualifying deferral fraction q, its annual value in $ millions equals G times q times (70 + 40), with $70/kW-year transmission and $40/kW-year distribution used as separate reference inputs. The 25%, 50% and 100% curves are assumptions about genuinely deferred infrastructure and not empirical estimates. The transmission input is a rounded ERCOT cost allocation rather than

marginal construction cost, and the distribution input derives from a study of residential solar, not transmission-connected large loads. Genuine avoided costs could be zero unless local planning studies show identified upgrades can be delayed or avoided with enforceable performance. Neither this potential benefit nor reallocated 4CP charges can automatically be subtracted from the ten-month wholesale measure. [28,29]



Appendix Figure A2. Conditional annual T&D investment-deferral sensitivity. The 25%, 50% and 100% values are assumptions about genuinely avoidable capacity, not estimates of actual ERCOT construction avoided. The $70/kW-year transmission reference comes from cost allocation, and the $40/kW-year distribution reference from a residential-solar literature proxy. Do not net this figure against the annual 4CP allocation estimate or the ten-month wholesale repricing result. [28,29]

A10. CHR and possible compensation arrangements

The Compute Heat Rate Index provides a quarterly reference for the electricity-price tolerance of a defined mix of AI inference workloads. The Q3 2026 blended long-run CHR is $5,631.40/MWh and the dispatch-focused CHR-D is $9,070.82/MWh. Training is excluded from the Q3 blended Index. CHR considers full-cost economics, while CHR-D uses avoidable costs relevant to curtailing existing computation. These are reference thresholds rather than forecasts of facility profitability, invoices to data centers or prescribed levels of ratepayer compensation. Neither value is an input to the ERCOT repricing calculation. [3,17]

An actual ratepayer-protection policy would need an evidence-based view of incremental customer costs and appropriate sharing among customers and developers. A separate, optional CHR-referenced community fund could base a negotiated fraction on the published National CHR Index times independently metered MWh, settled quarterly. That would require public governance, caps, anti-gaming controls, treatment of behind-the-meter structures and careful legal and tax review. This paper does not calculate a warranted rate for either mechanism. [3,11,17]

The wholesale-impact measure divides modeled gross repricing by 424,768,667 MWh of historical ERCOT consumption in the 86,076 matched pricing intervals. The hypothetical full-offset charge instead divides the same repricing numerator by incremental demand actually served under each operating strategy in those same priced intervals. The second denominator therefore reflects interval-specific curtailment, rather than the full-study capacity factors. The

first measure uses the gross-dollar totals rounded to $0.001 billion in the ERCOT report; the second uses the ERCOT interval-matched compensation calculation. Neither is an actual tariff, a retail bill estimate or a calculation of net system benefits. The 7 to 10 GW cases remain subject to the disclosed offer-stack boundary limitation. [4-6]

A10.1. Per-MWh results for all modeled load levels

Appendix Table A6. Modeled gross wholesale repricing per MWh of existing ERCOT consumption, in $/MWh. [4-6]

Added load	No curtailment	Peak-neutral	$150/MWh rule
1 GW	$2.10	$2.10	$2.06
2 GW	$4.03	$4.03	$3.96
3 GW	$5.85	$5.83	$5.71
4 GW	$7.88	$7.84	$7.60
5 GW	$10.45	$10.35	$9.90
6 GW	$14.85	$14.54	$13.58
7 GW	$36.13	$33.06	$20.15
8 GW	$98.29	$85.86	$28.89
9 GW	$213.24	$180.06	$37.14
10 GW	$339.60	$278.29	$44.32

Appendix Table A7. Illustrative additional cost per MWh of new electricity served in order to fully offset wholesale repricing, in $/MWh. [4-6]

Added load	No curtailment	Peak-neutral	$150/MWh rule
1 GW	$124.27	$124.28	$122.63
2 GW	$119.42	$119.51	$117.65
3 GW	$115.50	$115.66	$113.47
4 GW	$116.65	$116.80	$113.65
5 GW	$123.78	$123.66	$118.78
6 GW	$146.60	$145.29	$136.28
7 GW	$305.67	$284.11	$174.20
8 GW	$727.53	$647.87	$220.37
9 GW	$1,403.05	$1,212.20	$255.23
10 GW	$2,011.05	$1,693.53	$278.80

Notes: Model period December 5, 2025 through October 4, 2026; all ten load levels use the same 86,076 priced intervals. Table A6 is based on $ billions rounded to three decimals in the ERCOT summary; Table A7 uses the scenario-specific priced-interval served energy. The $150/MWh response is a curtailment trigger, not a limit on a hypothetical full-offset charge. [4-6]

A11. Funding, independence and next analysis

The ERCOT wholesale-market research was developed and funded by RP Energy, Inc., not commissioned by a utility, data-center developer, generator or advocacy organization. The conclusions are those of the author and RP Energy, Inc. Hans Royal, a co-founder of RP Energy, created the Compute Heat Rate. RP Energy has an interest in CHR's institutional adoption and potential commercial licensing. The CHR-referenced community fund is a possible future application discussed in this paper.

The priority follow-ups are location-specific ERCOT load-zone and nodal price incidence; actual settlement-based transmission and distribution cost allocation and identified deferrable projects; enforceable structured flexibility and battery/virtual power plan operation under realistic duration and notice; additional low-cost generation by hour and location; and customer exposure after hedging and retail contracts. PJM energy and capacity-market analysis and a regulated-utility fuel-cost study would test how far the framework carries beyond ERCOT.



Appendix Figure A3. Coverage of ten-month ERCOT observations. Of 87,530 load intervals, 86,076 are usable for modeled wholesale repricing. Five days lack usable supply curves, and 14 other five-minute intervals lack matching System Lambda. [4-6]

References and primary data sources

[1] Norris, Tyler H., Timothy Profeta, Dalia Patino-Echeverri and Adam Cowie-Haskell. Rethinking Load Growth: Assessing the Potential for Integration of Large Flexible Loads in US Power Systems. Duke Nicholas Institute, February 2025. https://nicholasinstitute.duke.edu/publications/rethinking-load-growth

[2] Levitt, Andrew, Johannes P. Pfeifenberger, Joe DeLosa III and Nathan Felmus. Solutions for Timely Interconnection of Large Loads: Enabling Non-Firm Transmission Service. The Brattle Group, prepared for Google, September 2026. https://www.brattle.com/insights-events/publications/new-brattle-report-explores-solutions-for-timely-interconnection-of-large-loads/

[3] Royal, Hans. The Compute Heat Rate: Quantifying AI-Driven Electricity Price Tolerance and Its Implications for Wholesale Market Repricing. February 28, 2026, revised June 4, 2026. SSRN 6322318. http://dx.doi.org/10.2139/ssrn.6322318

[4] ERCOT. NP3-908-ER, 2-Day SCED Energy Curves. Report Type ID 13054. https://www.ercot.com/mp/data-products

[5] ERCOT. NP3-910-ER, 2-Day Real Time Generation and Load Data Reports. Report Type ID 13056. https://www.ercot.com/mp/data-products

[6] ERCOT. NP6-322-CD, SCED System Lambda. Report Type ID 13114. https://www.ercot.com/mp/data-products

[7] ERCOT. Historical Load Data Archives. https://www.ercot.com/gridinfo/load/load_hist

[8] Google. Why We're Backing America's Existing Nuclear Plants. October 6, 2026. https://blog.google/company-news/why-were-backing-americas-existing-nuclear-plants/

[9] Constellation Energy and Google. Landmark Agreement to Bring 890 MW of New Nuclear Capacity to PJM Grid. October 6, 2026. https://www.constellationenergy.com/news/2026/10/google-and-constellation-announce-landmark-agreement-to-bring-890-mw-of-new-nuclear-capacity-to-pjm-grid.html

[10] Google. Scaling Clean Energy and Reliability in Michigan. March 17, 2026. https://blog.google/innovation-and-ai/infrastructure-and-cloud/global-network/clean-energy-reliability-michigan/

[11] Alaska Permanent Fund Corporation. Permanent Fund structure and Percent of Market Value methodology. https://apfc.org/the-fund/fund-structure/

[12] U.S. Environmental Protection Agency. 2025 GHG Emission Factors Hub (Jan. 2025), Table 6 (eGRID). https://www.epa.gov/system/files/documents/2025-01/ghg-emission-factors-hub-2025.pdf

[13] U.S. Environmental Protection Agency. Avoided Emissions and Generation Tool (AVERT). https://www.epa.gov/avert

[14] ERCOT. Energy-Only Wholesale Market. ERCOT Grid Insights, October 16, 2025. https://www.ercot.com/files/docs/2025/10/16/ERCOT-Grid-Insights-Energy-Only-Wholesale-Market.pdf

[15] PJM Interconnection. Capacity Market, Reliability Pricing Model. https://www.pjm.com/markets-and-operations/rpm

[16] PJM Interconnection. Energy Market. https://www.pjm.com/markets-and-operations/energy

[17] Royal, Hans. CHR Reference Values, Q3 2026. Compute Heat Rate Index, September 1, 2026. https://computeheatrate.com/chr-index

[18] McEntush, Ryan. Why Texas Is Making Data Centers Wait. Andreessen Horowitz, October 7, 2026. https://a16z.com/why-texas-is-making-data-centers-wait/

[19] Wiser, Ryan et al. Retail Electricity Price Trends and Drivers: Data Update, 2026 Edition, with July 2026 update. Lawrence Berkeley National Laboratory and The Brattle Group. https://eta-publications.lbl.gov/publications/retail-electricity-price-trends-0

[20] Brancucci, Carlo, Dylan Cutler and Jesse Jenkins. Flexible Data Centers: A Faster, More Affordable Path to Power. Camus, encoord and Princeton University ZERO Lab, December 2025. https://www.camus.energy/flexible-data-center-report

[21] Korgel, Brian A., Shipra Mishra and Tiffany Wu. Powering Texas at the Speed of Business: Proceedings from Phase 3 of the Large Load Symposium at The University of Texas at Austin. The Energy Institute at The University of Texas at Austin, March 23, 2026. https://doi.org/10.26153/tsw/63129

[22] U.S. Department of Energy, Office of Electricity. Strengthening America's Grid Supply Chain. August 10, 2026. https://www.energy.gov/oe/articles/strengthening-americas-grid-supply-chain

[23] McKenna, Killian, Sherin Ann Abraham and Wenbo Wang. Major Drivers of Long-Term Distribution Transformer Demand. NREL/TP-6A40-87653, 2024. https://doi.org/10.2172/2309697

[24] ERCOT. RTC+B Implementation Go-Live Complete for Operating Day December 5, 2025. Market Notice M-F110525-04, December 5, 2025. https://www.ercot.com/services/comm/mkt_notices/M-F110525-04

[25] The White House. Ratepayer Protection Pledge. March 4, 2026. https://www.whitehouse.gov/releases/2026/03/ratepayer-protection-pledge/

[26] U.S. Environmental Protection Agency. Emergency stationary engine operating restrictions. 40 CFR 60.4211(f) (Subpart IIII, stationary compression-ignition engines) and 63.6640(f). https://www.ecfr.gov/current/title-40/chapter-I/subchapter-C/part-60/subpart-IIII/subject-group-ECFR7cc160b57a360ac/section-60.4211

[27] Michigan Public Service Commission. MPSC Approves DTE Electric Energy Contracts for Google Data Center, Protects Other Customers. October 1, 2026. https://www.michigan.gov/mpsc/commission/news-releases/2026/10/01/mpsc-approves-dte-electric-google-data-center-and-contracts-protects-customers-from-project-cost

[28] ERCOT Board of Directors. Strategic Discussion on Resource Adequacy and the Role of Demand Response, April 20-21, 2026. Presentation cites a 2025 ERCOT 4CP transmission charge of $68.55/kW-year in a Charles River Associates discussion. https://www.ercot.com/files/docs/2026/04/13/11.1-Strategic-Discussion-on-Resource-Adequacy-and-the-Role-of-Demand-Response.pdf

[29] Dunsky Energy + Climate Advisors. Value of Residential Solar in Texas, prepared for Texas Solar Energy Society, July 16, 2024. Figure 3-13 presents a cross-study average avoided distribution capacity benchmark of approximately $40/kW-year (2024 dollars), with location-specific qualifications. https://txses.org/wp-content/uploads/2024/07/Value-of-Residential-Solar-in-Texas.pdf

[30] PJM Interconnection. Powering Reliability Through Market Design: Addressing Rising Demand and Constrained Supply, and Stimulating Investment To Support Durable Reliability. May 6, 2026, pp. 49·50. https://www.pjm.com/-/media/DotCom/library/reports-notices/special-reports/2026/20260506-powering-reliability-through-market-design.pdf

[31] Google. We are working with Voltus to build a 1-gigawatt virtual power plant. October 22, 2026. https://blog.google/innovation-and-ai/infrastructure-and-cloud/global-network/virtual-power-plant-1-gigawatt/

[32] Voltus. Google and Voltus Announce One Gigawatt Bring Your Own Capacity Agreement in PJM. October 22, 2026. https://www.voltus.co/press/google-voltus-1gw

