DAVID S. LAPP
PEOPLE'S COUNSEL

WILLIAM F. FIELDS
DEPUTY PEOPLE'S COUNSEL

JULIANA BELL
DEPUTY PEOPLE'S COUNSEL

OPC

OFFICE OF PEOPLE'S COUNSEL

State of Maryland

6 ST. PAUL STREET, SUITE 2102
BALTIMORE, MARYLAND 21202
WWW.OPC.MARYLAND.GOV

BRANDI NIELAND
DIRECTOR, CONSUMER ASSISTANCE UNIT

CARISSA RALBOVSKY
CHIEF OPERATING OFFICER

October 14, 2025

Mr. David Mills, Chair, PJM Board of Managers
Mr. Manu Asthana, President and CEO
PJM Interconnection, LLC
2750 Monroe Boulevard
Audubon, Pennsylvania 19403

Re: Large Load Adjustments for the 2026 Load Forecast

Dear Mr. Mills and Mr. Asthana:

The Maryland Office of People's Counsel (MPC) respectfully submits this letter to the PJM Board of Managers requesting its intervention to address the extraordinary risks to existing customers that the data center "land rush" presents. These challenges need to be addressed immediately. The current forecasting cycle for 2026 will set the parameters for upcoming capacity auctions as well as impact transmission planning and procurement, resulting in billions of dollars in costs for PJM customers.

Recent new and additional risks to customers' costs have been triggered by gargantuan large-load adjustment (LLA) requests of transmission owners and load-serving entities far above previous requests. We compiled the 2026 LLA requests for increases in "demand" and "capacity"[1] to illustrate how they comprise significant increases over the 2025 PJM forecast, which itself already incorporated large increases from current load levels. By our calculation, using the more conservative "demand" figures, the aggregated increases go from about 19 GW for 2028 in the 2025 forecast to an increase of 32 GW in 2028 (a 68 percent increase); and from an increase of about 33 GW in 2030 in the PJM 2025 forecasted LLAs, to approximately 60 GW in 2030 (an 83 percent increase). In terms of non-coincident PJM footprint wide peak load, the LLA requests in aggregate represent a 20 percent increase in 2028 over the 2025 peak (or a 32 GW increase), and, in 2030, a 37 percent increase over the 2025 peak load (60 GW).

[1] Citations and references for statements in this letter are provided in the attached comments that have been submitted to the Load Analysis Subcommittee.

To put these proposed increases in perspective, the 2030 forecasted load *increase* from this summer's PJM peak load is larger than the peak load currently served by the California Independent System Operator, comprising the majority of electric load in the state of California, and by 2032 the load increase postulated for PJM is close to that currently served by the Electric Reliability Council of Texas.

The attached comments detail how the massive PJM-wide increases are built on the shaky foundation of widely varying standards that PJM transmission owners use to include data center projects in their requests. Even for the transmission owners with more rigorous approaches, stakeholders do not know critical details about the thresholds that transmission owners are using to incorporate data centers into their forecasts.

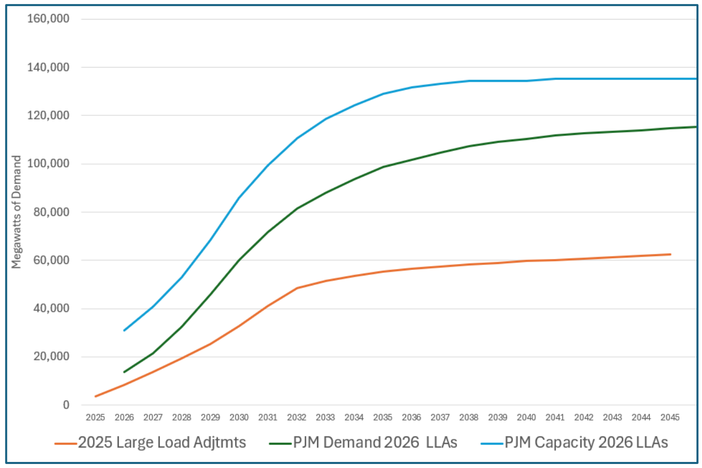

Preliminary 2026 large load adjustment requests (demand and capacity) v. PJM 2025 load report

Beyond those varying standards, we have submitted the attached comments and recommendations for consideration by the Load Analysis Subcommittee in its current work on the forecasts that will be incorporated into the PJM 2026 annual report. These issues must be addressed before the 2026 load forecast is adopted.

- The aggregated LLAs massively overshoot recent national level forecasts of data center development over the next five years, equating PJM's anticipated data center load growth with estimates for the entire country.

- The aggregated LLAs are wholly disproportionate to the levels of data center investment that would be required to support the forecasts. Estimates of the capital investment required to build a 1 GW data center are about $30 billion. New data center load of 32 GW—the aggregated LLAs for 2028—would require more than *$1 trillion* in capital expenditure, all in PJM.

- Estimates of the total amounts of annual U.S. capex expenditure for data center infrastructure in 2024 and projections for 2025 fall far short of the levels of capex that would be necessary to fulfill the LLA requests in PJM.

- Recent trends in AI and data center deployment suggest that geographically concentrated deployment is becoming less likely and necessary. If that trend is realized, it could compromise the focused buildout of the data center cluster in Northern Virginia and Central Ohio that drive much of the LLA load requests.

- Recent announcements of huge data center infrastructure developments in the Western U.S. may undercut the LLAs' estimate. These developments, anchored by anticipated economies of scale and efficiencies in land assembly as well as increasing latency tolerances for certain types of AI development, portend a competitive alternative to the buildout of data centers within the PJM footprint.

- Data center physical infrastructure for computing and AI applications is under continuous and rapid change due to technological innovation—with important implications for data center energy use. These changes are likely to be accompanied by improvements in energy efficiency and load flexibility

Given these reasons for skepticism about scale of development, the LLA requests create a serious risk of huge amounts of stranded costs, both for capacity and transmission assets. If the huge financial risks currently propelling the data center sector lead to major losses within the sector, existing, captive customers will be at significant risk of bearing these stranded costs. Accordingly, the electric industry, as it plans to serve this new kind of customer, should take a measured and prudent view, meaning much more scrutiny than that being accorded presently of important intermediate and long-term trends propelling the data center sector.

Outside of stranded-costs risks, the proposed LLA adjustments will impose massive costs on customers across PJM if included in the 2026 load forecast. They will drive capacity market costs for the auction scheduled for next June as well as the need for additional transmission projects that will burden customers and communities. In order to minimize these risks for customers, PJM must:

- develop a rigorous, fully transparent and uniform process across all zones in PJM to review the LLA requests;
- establish specific requirements for inclusion in the load forecast, including required contractual terms and financial requirements; and
- specify project milestones—such as site control, the ability to procure sufficient hardware for the project, and the ability of the local utility

to complete the necessary projects—before allowing the data center to be connected to the system.

These review criteria should be fully transparent to all stakeholders.

Billions of dollars in costs and tremendous disruption from transmission development for the residents and businesses in the PJM region are at stake in determining the 2026 load forecast. MPC requests that the Board carefully review the attached comments to the LAS as well as the views of other stakeholders and, prior to finalizing the 2026 load forecast, develop a transparent review process for the LLA requests that results in a forecast that is not at odds with the available information on data center development.

Very truly yours,

David S. Lapp
People's Counsel

Enclosure: Comments to Load Analysis Subcommittee

Maryland Office of People's Counsel ("MPC")
Comments regarding the current review cycle for preparation
of the PJM 2026 annual load forecast report.

Introduction and Summary

MPC prepared these comments to inform PJM's current and pending review of the large load adjustment ("LLA") filings by participating transmission owners ("TOs") and load serving entities ("LSEs") operating in PJM's footprint. These comments address PJM's current load forecasting cycle, intended to result in PJM's 2026 annual load report, anticipated to be published in early 2026.

Similar to last year's annual forecasting cycle culminating in the 2025 PJM annual load report, this year's LLA requests include huge additive increases in forecasted LLAs looking out next year and over the next decade. The huge increase in forecasted future loads, primarily if not solely from new data centers, has the attributes of the 1889 Oklahoma land rush, but this time for electric grid access.[1]

This land rush to stake out a claim to access the electric system arises, in part, from the huge financial bets in the face of risks placed by and currently driving the data center infrastructure sector in pursuit of its strategic need for scarce resources. The land rush is further seemingly enabled by the incentives of the TOs submitting the LLA requests to build transmission facilities and expanding their regulated rate base, where captive ratepayers bear risks from electric facility overbuild and stranded costs if the loads do not materialize.

The PJM annual load forecast and its prospective incorporation of the TOs/LSEs' LLAs is the field where these forces are playing out. Much greater rigor and review of these LLA requests must occur in the current load forecast cycle to counterbalance these major risks and incentives—to mitigate against possible huge amounts of stranded assets and misallocation of capital investment and corresponding costs and risks for customers.

We compiled the aggregate of the TOs' and LSEs' individual LLA requests for "demand" and "capacity" increases recently posted by and then compiled together by PJM.[2] The LLAs currently under review in aggregate go from an increase from today's

---

1 The Oklahoma land rush started at high noon on April 22, 1889, when an estimated 50,000 people lined up at the start to lay stake to a piece of about two million acres of Native American territory. *Land Rush of 1889*, https://en.wikipedia.org/wiki/Land_Rush_of_1889 (visited Oct. 9, 2025).

[2] PJM separately compiled the TOs' load requests for "capacity" with much higher values. By way of example, the aggregate "capacity" requests are an incremental 31 GW in 2026 rising to 86 GW in 2030

maximum summer 50/50 demand of about 19 GW increase for 2028 in the PJM 2025 forecast, based on the 2025 forecast's LLAs (from Table B-9 of the PJM 2025 forecast) to an increase of 32 GW in 2028 (a 68 percent increase); and from an increase of about 33GW in 2030 in the PJM 2025 forecast, to approximately 60 GW in 2030 (an 83 percent increase). In terms of non-coincident PJM footprint wide peak load, the pending LLA requests in aggregate represent a 20 percent increase in 2028 over 2025 (32 GW increase), and, in 2030, a 37 percent increase (60 GW increase).

To put these increases in perspective, the 2030 forecasted load increase comprised by the current LLA filings in aggregate, from this summer's PJM peak load, is larger than the peak load currently served by the California Independent System Operator ("CAISO"), comprising the majority of electric load in the state of California; by 2032 the load increase postulated for PJM is close to that currently served by the Electric Reliability Council of Texas ("ERCOT"), serving the majority of electric load in the state of Texas.

A graphical depiction of these changes is set forth below:

Figure 1: Preliminary 2026 large load adjustment requests (demand and capacity) v. PJM 2025 load report

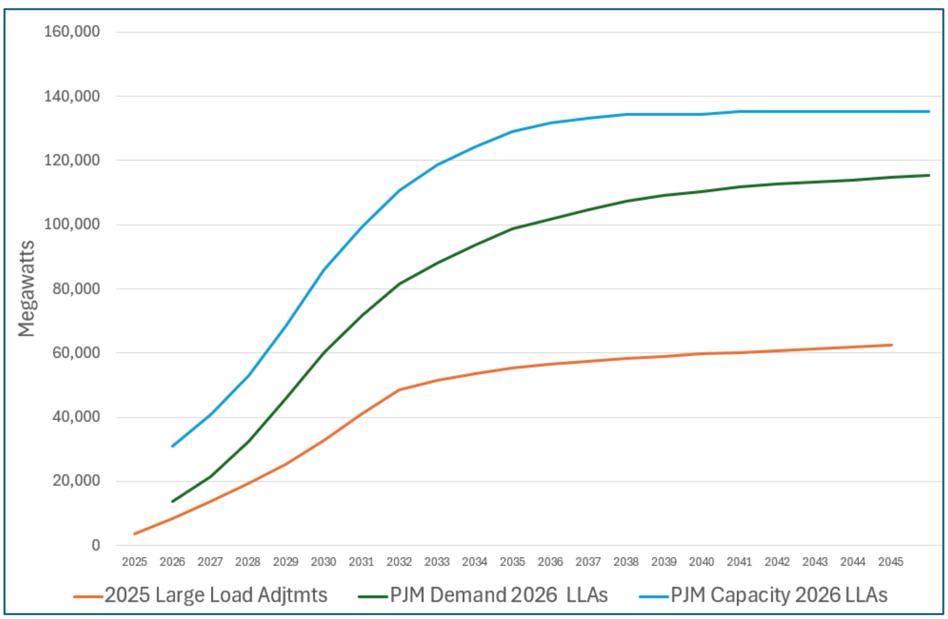

---

vs. 14 GW (2026) and approx. 60 GW (2030) incremental LLA requests for "demand." For present purposes, we conservatively assume that the "demand" contained in the LLA requests (rather than the "capacity" LLAs) are the values under review by PJM for possible inclusion in the load forecast.

The following table depicts these changes calculating the relatively more "conservative" LLA requests for "demand" only (rather than for the much higher "capacity" LLAs) compiled by PJM.

Table 1: Preliminarily 2026 Large Load Adjustments (Demand)

	Year	2027	2028	2029	2030	2031	2032
Summed Prelim PJM 2026 Large Load Adjustment Requests (LLA) for "Demand"	MWs	21,402	32,391	45,796	60,117	71,552	81,463
Total LLAs in 2025 PJM Forecast - Table B9	MWs	13,668	19,320	25,334	32,795	41,101	48,531
Prelim 2026 LLAs less 2025 LLAs	**MWs**	**7734**	**13071**	**20,462**	**27,322**	**30,451**	**32,932**
2026 LLAs % over 2025 LLAs	**%**	**56.6%**	**67.7%**	**80.8%**	**83%**	**74.1%**	**67.9%**
Cumulative % LLAs 2026 Increase over 2025 NCP Peak	%	13%	20%	28%	37%	44%	51%
NCP RTO Summer 50/50 Demand (MWs) Table B1	160,961						

If these forecasts are included in PJM's 2026 annual load report, they will anchor and scale future PJM capacity market procurements and transmission expansions—causing a massive and unprecedented expansion in the electric supply grid and increased cost responsibility of PJM's electric consumers—measured in many billions of dollars, if even feasible. Thus, they put even more stress on PJM's Critical Issues Fast Path ("CIFP") process currently in process to address the supply scarcity that the PJM load forecast predicts in all cases even before this recent round of LLA filings of the PJM TOs/LSEs.

Meeting this singular moment requires a more robust review during this forecast cycle to avoid or mitigate the potential for massive overbuilding and mis-allocation of capital. Given the massive increases in the LLA filings, there is substantial room that the forecasts will not be realized, causing significant stranded costs. This risk exists even recognizing that PJM has sought to do a better review of LLAs during this year's forecast cycle, when compared with the process conducted last year.

At a minimum, great skepticism is warranted about the individual TOs' LLA filings:

- The aggregate of the individual TOs' LLAs, driven almost entirely by forecasts of data center load growth, massively overshoot recent national level forecasts of data center development over the next five years, in many cases equating PJM's anticipated load growth due to data centers with that projected for the entire country, when the PJM footprint only covers approximately one-fifth of the US in population and similar shares of other measures of economic activity.[3]

- The aggregate level of load comprised by the LLAs' requests is wholly disproportionate if extrapolated to the likely data center investment required to support the levels of these forecasts. Estimates of the capital investment required to build a 1 GW data center are in the neighborhood of $30 billion.[4] New data center load of 35 GW—the aggregated LLAs' projection in 2028 for PJM—would require more than *$1 trillion* in capital expenditure, all in PJM.

- In 2024, the aggregate annual capex expenditure in the U.S. for AI data center infrastructure by the top four AI hyperscalers, Tier 2 Cloud operators and enterprises and government infrastructure is estimated at $212 billion.[5] In 2025, the top four AI hyperscalers

[3] *See, e.g.*, the Department of Energy's ("DOE's") recent report on electric reliability estimates 50 GW of data center load increases by 2030 across the entire US, and about 15 GW of that load occurring within the PJM footprint. DOE, *Resource Adequacy Report: Evaluating the Reliability and Security of the United States Electric Grid* (July 2025) (p. 16, national estimate of 50 GW in load growth by 2030; PJM allocation of 29% of total or about 15 GW, pp. 17 and 28); EPRI, *Scaling Intelligence: the Exponential Growth of AI's Power Needs* (2025 White Paper) (projecting 50 GW in "AI Demand" in the US for both training and inference by 2030, p. 22); Kou and Limandibhratha, *Power for AI: Easier Said Than Built*, Bloomberg NEF (April 15, 2025) (Total US data centers' load growing from 35 GWs (2024) to 78 GWs (2035) (or a 45 GW increment by 2035); Goldman Sachs, *Generational Growth: AI, data centers and the coming US power demand surge* (April 28, 2024) (estimates 47 GW of incremental demand in the US driven by data centers through 2030; PJM share is 15% of the data center load (p. 21) (or about 7.5 GW).

[4] Tabby Kinder, *Absolutely Immense: the companies on the hook for the $3Trn AI building boom*, Financial Times (Aug. 14, 2025) (quoting an executive from Blackstone); Bye, *Can We Build a Five Gigawatt Data Center?* Asterisk Magazine (2025).

[5] See Fermi LLC, Form S-11 (Aug. 6, 2025). (draft registration of initial public offering prospectus), p. 11 (Fermi LLC is developing an up to 15 GW data center campus in Amarillo, Texas). *See also*, JP

(Amazon, Microsoft, Google, Meta) and OpenAI plan are reported to be on track to spend in the U.S. $325 billion on data center infrastructure.[6] The PJM footprint is home to approximately 20% of U.S. population and other measures of economic activity; it appears inconceivable that PJM would see the investment levels necessary to meet the 2026 LLA requests that would require more investment in PJM than current estimates of investments taking place nationally.

- Recent trends in AI and data center deployment suggest that geographically concentrated data center deployment is becoming less likely and necessary. If so, there will be broader geographic deployment of data center infrastructure across the U.S. and globally, rather than the focused build-out of the data center clusters in Northern Virginia and Central Ohio, which comprise a large part of PJM's LLA load requests.[7]

- Recent announcements of huge data center infrastructure developments almost all in the Western U.S., not included in the national estimates cited previously, undercut the PJM TOs/LSEs' LLAs aggregate estimate of data center development within the PJM footprint.[8] These developments, anchored by anticipated

---

Morgan Asset Management, *Market Insights, Guide to the Markets, US 4Q 2025 as of September 2025*, p. 22 (showing $211 Billion in 2024 in capex for Alphabet, Amazon, Meta, Microsoft and Oracle).

[6] Metz and Weise, *What exactly are AI firms attempting to build?* New York Times, pp. B1, B6 (Sep. 29, 2025). See also, JP Morgan (Sep. 30, 2025) (projecting $347 Billion in capex from Alphabet, Amazon, Meta, Microsoft and Oracle in 2025).

[7] Fermi LLC, Form S-11, draft registration statement (Aug. 6, 2025), pp. 11 ("Gravitational Pull": Latency is not a primary constraint for many large-scale GenAI workloads, particularly AI model training. As a result, hyperscalers are prioritizing access to power and speed-to-market over proximity to traditional Tier 1 data center markets. This shift coincides with severe power constraints in legacy Tier 1 markets such as northern Virginia, where utility moratoriums and lack of available power and infrastructure have stymied new data center development. Accordingly, new data center builds are rapidly expanding into previously Tier 2 and Tier 3 markets which are rapidly emerging as new centers of compute infrastructure activity due to available grid capacity. This trend reflects a broader decentralization of compute infrastructure in the AI era, where access to grid-independent, scalable, massive amounts of power is driving hyperscalers' decisions.").

[8] OpenAI, Press Release (Sep. 23, 2025), *OpenAI, Oracle and Softbank expand Stargate with five new AI data centers* (10 GW commitment for data center developments, $500 billion investment for

economies of scale and efficiencies in land assembly as well as increasing latency tolerances for certain types of AI development, portend very large data center capacity in locations other than in PJM, competitive at least in part with developments within the PJM footprint, and an alternative to the build-out of data center infrastructure capacity within the PJM footprint.

- The physical infrastructure deployed at data centers to support computing and AI applications is under continuous and rapid change. Technological innovation has significant implications for data center energy use, including improvements in energy efficiency and load flexibility. They include an increasingly faster cadence of new generations of chip technology, likely bringing changes in supporting technology, such as different methods of cooling and other "balance of plant" equipment required to make data centers functional. Other areas of rapid innovation include software code, changes in microchip architecture, reconfiguration of the conduct of AI applications, both geographically and as affecting energy usage, reconfiguration of equipment usage to support AI applications, and changes in the balance between "edge" and data center hosting of computational activity.[9]

  Power requirements also vary by type of data center. Colocated (or "colos") or hyperscaler direct control and developed data centers may have very different electric usage patterns and ability to control and manage electric usage within an individual data center site. Further distinctions, affecting energy usage, may apply to "turn-key" and "powered shell" developments.[10] PJM's current process

---

facilities in Abilene, Shackelford County and Milam County, Texas and Dona Ana County, New Mexico, Lordstown, Ohio, and an undisclosed site in the Midwest); Fermi LLC, Initial Public Offering prospectus (11 GW data center campus under development in Amarillo, Texas). SEC, Amendment No. 1 to Form S-11 Registration Statement (Sep. 24, 2025), p.1.

[9] The literature is vast regarding the pacing and trends in innovation in computer processing technology. *See, e.g*., Stanford University, *The Stanford Emerging Technology Review 2025* (2025), chaps. 1 and 8.;Rahman, R. *Leading ML Hardware Becomes 40% More Energy Efficient Each Year*, Epoch AI (2024); Nadel, *Opportunities to Use Energy Efficiency and Demand Flexibility to Reduce Data Center Energy Use and Peak Demand*, ACEEE (Oct. 2025).

[10] *See, e.g*, Urban Land Institute, *Local Guidelines for Data Center Development* (2024), pp. 7-8 (distinguishing among enterprise, colocation (retail), wholesale, telecom and hyperscaler data centers); Equinix, *What is a Data Center? What are different types of data centers?* (2022, updated to 2024);

does not distinguish or document the nature of the large load interconnection being requested. As a result, it is insensitive to the heterogeneous requirements of each site regarding power costs, latency, reliability and other factors.

PJM should account for each of these factors, which individually and collectively portend significant changes in PJM forecasts of data center energy use. They could be analyzed through probabilistic modelling used for different applications, such as the BRA and RTEP. Forecasts should be discounted to reflect probable uncertainties. In any event, the LAS should address these uncertainties directly and transparently.

These highlights are explored in detail in three sections below.

Section 1 addresses the continuing need for complete, uniform, robust and comparable documentation and support analysis across the filings of the TOs/LSEs. This completeness, uniformity, robustness, and documentary support is not apparent across the TO/LSE LLAs filings.

Section 2 addresses the question about how the current CIFP (and the putative NCBL (non-capacity backed load) proposal, now withdrawn, or its possible replacement) can and likely would affect the load forecast.

Section 3 addresses the critical need to expand the due diligence surrounding the load forecast, to include the data center sector's financial and risk circumstances and supply chain issues. PJM's current load forecasting process largely ignores those circumstances. While not normally part of electric load forecasting, the financial and risk circumstances warrant PJM's attention now given the massive costs and expenditures facing customers and the associated risks of stranded costs accompanying those costs and spending.

Discussion

Section 1: PJM should fix its load forecasting to promote uniformity, completeness, documentation.

Accurate forecasting of data center electric load growth is a major challenge for electric utilities. Challenges result from the lack of prior experience, the "claim staking" to grid

---

Prince William County, *2024 Data Center Industry Tax Revenue Report* (2025) (distinguishing between "turn-key" and "powered shell" data centers for purposes of property tax assessment. Power capacity is used in valuation for property tax purposes only for turnkey developments.).

access anchored by customer service requests to the electric utilities, and rapid pace of technological developments.[11] Best practices in load forecasting of data center load is still in its infancy and subject to significant uncertainties, larger than the inherent uncertainties faced by forecasting generally.[12]

1.1—The land rush

This staking of claims to access the electric grid—the "land-rush" mentioned above—describes recent experience. In the PJM footprint, projections of rapid future load growth due to data centers started with the Dominion service area, spread to AEP in last year's forecast cycle, and now with this most recent filing of LLAs includes PPL and ComEd, as well as several of the First Energy affiliates. In this environment, many developers in the data center ecosystem file electric service requests for multiple sites across multiple electric utilities to offset the risk that any single location might be foreclosed because of siting, land use, electric grid limitations or other obstacles that surface during the data center development process. The claim staking is exacerbated by efforts to find the most accommodating local or state governments willing to partner, reduce risk, and provide subsidies, as well as competition among data center developers to secure the same or nearby sites or overlapping points of grid access.

The data center development ecosystem also is changing, with the entry of many more not fully integrated participants besides the hyperscalers, including often highly leveraged neoclouds,[13] adding to the pile-on in electric service requests seeking to lay claim to headroom on the electric grid. As relative scarcity of available suitable sites goes up and preferred grid locations go down due to earlier "claims" (on sites, but also on the grid through earlier in time electric service requests) that preempt later entrants, developers can also create "moats" against competing developers.

[11] See, e.g., Goldsmith and Byrum, *Powering the US Data Center Boom: Why Forecasting Can Be So Tricky*, World Resources Institute (Sep. 17, 2025); Frick and Srinivasan, *Large Load Literature Review* (Aug. 2025 Update), Energy Technologies Area, Berkeley Lab; NERC, *Characteristics and Risks of Emerging Large Loads, Large Loads Task Force White Paper* (July 2025); Koomey, et al, *Electricity Demand Growth and Data Centers: A Guide for the Perplexed* (Feb. 2025), Bipartisan Policy Center/Koomey Analytics; Wilson et al., *Strategic Industries Surging: Driving US Power Demand* (Dec. 2024) Grid Strategies.

[12] *Id.*

[13] See discussion below in the Appendix.

1.2 The need for best practices across PJM

In light of these challenges, a developing best practice is for electric utilities to require an ascending set of contractual commitments from the data center developer, paced over the data center development life cycle, to reflect:

(i) greater "skin in the game" by the data center developer and, therefore, assurance of a higher likelihood of data center load showing up commensurate with a greater share of the costs of serving that load born by the developer; and

(ii) greater protections to the utility and its other customers from the risks that the data center load does not occur by allowing, through the contractual commitments with the data center developer, for the funding by the developer of the costs expended to expand the grid to serve the new data center load.

Absent these contractual commitments, data center load growth forecasts should be deeply discounted. Even with these contractual commitments, the realism of the electric load forecasts estimated, in the first instance, by data center customers should be scrutinized, as there is wide variation in the level of cost commitment and obligation that these contractual engagements actually entail. In the current speculative environment entailing grid expansion to serve very large incremental loads, having signed "agreements" alone, without specifying the level of commitment, is not probative of the realism of the estimated future loads.

Echoing and paralleling this trend, PJM in its guidance for load forecast adjustment development by its TOs and LSEs explicitly asked the TOs to provide information regarding the nature of the contractual obligations.[14] PJM's guidance is permissive, however. It does *not* mandate the TOs and LSEs to:

- follow a uniform basis for contracting with data centers,
- establish uniform and consistent disclosure across the TOs and LSEs of the specific levels of commitment for which data center customers are obligating themselves; or
- state the stage of the contracting process that the developer has completed.

[14] See PJM Manual 19: *Load Forecasting and Analysis, Attachment B: Load Forecast Adjustment Guidelines.*

In sum, without requiring this information to be provided on a mandatory basis, it is not possible to know what level of commitment the electric loads for such customers should be included in the TO's LLA for incorporation into PJM's annual load forecast.

1.3: Lack of consistency and best practices impacting 2026 preliminary large load adjustments

PJM's publicly posted documentation accompanying the current round of LLAs fails to inform the level of binding commitment from the data center customers, and it is not uniform across the filings by the TOs/LSEs.

a. *Dominion Energy*

Dominion's LLAs provide an example. Dominion has for several years anchored its forecast by ascending levels of contractual commitment from its data center customers, represented by engineering letters of authorization ("ELOAs"), construction letters of authorization ("CLOAs") and electric service agreements ("ECAs"). Each in the sequence of ELOAs, CLOAs and ECAs represent greater levels of contractual commitment by the customers signing these agreements.[15] The driver for the Dominion LLA request in the current PJM forecast cycle, however, apparently is the annual rate at which Dominion has historically been able to connect new large loads. Dominion's LLA submittal for periods after 2035 exceeds the demand represented by data centers with ECAs, which flat-line after 2028. Dominion's LLA request also fails to make clear what level of financial commitment is represented by the contractual commitments that are in place and whether and to what extent those commitments will cover the full cost of grid expansion to serve the incremental loads. Dominion does not identify transmission costs beyond the "last mile" or connecting substation directly serving the data center. It is also not clear whether the load increases are linked to Dominion's (or the regional grid's) ability to complete the transmission and distribution level expansion at or before the time the load increase is planned to occur.

Dominion is engaged in a proceeding before the Virginia State Corporation Commission ("VASCC") to establish a new "large load" retail tariff that would formalize a framework for consistent contractual engagement by and cost allocation to

[15] See Dominion Energy Virginia, *2025 20-Year Data Center Forecast, September 2025*. Item 4a to PJM LAS materials, Sep. 16, 2025.

the large load data centers, shifting some measure of cost responsibility to data centers for the grid expansion costs caused by the data center load additions.[16]

Notwithstanding the pendency of this proceeding before the VASCC and the proposal of the new tariff, Dominion has not adjusted materially its forecast of increases for large load this year from last year's forecast, when the retail tariff was not pending. When asked during the September 2025, PJM LAS committee meeting about this possible discrepancy, the Dominion representative stated that the pending retail tariff largely formalized its on-going practice and that other peer utilities were also establishing similar tariffs so that there was a low or no possibility of migration of the load away from Dominion's territory to other utility territories with less rigorous requirements that would likely have a material impact on Dominion's load forecasting.

It is not clear that the retail tariff as finally proposed, in fact, embeds Dominion's prior practice. It is also not clear whether and how Dominion will change its practices conforming to the new tariff with respect to earlier in time electric service requests. In addition, other than perhaps AEP and a few others, "peer" utilities have not, at least yet, adopted practices that are comparable to those claimed by Dominion.

In contrast to Dominion—and contrary to Dominion's rationale for its lack of adjustments to its LLA in the wake of the VASCC new tariff proposal—the introduction of AEP's retail large load tariff seems to have had a big impact in reducing large load interconnection requests.

b. *American Electric Power—Ohio*

AEP's recently adopted large load tariff, applicable to its Ohio service territory, requires increased financial commitments from data centers to address the need to put the data centers' "skin in the game" and protection of existing ratepayer interests (which the Dominion tariff is directed at as well).[17] It is not clear from the materials posted by PJM regarding the AEP LLAs how the AEP large load tariff varies from that utilized by Dominion. Also unclear is the level of financial commitment to the full cost of grid expansion required from data center customers or whether the commitments are paced to when AEP can actually construct the facilities to provide service at the levels requested. Moreover, after implementation of its new large load/data center tariff,

---

[16] VA SCC, Case No. PUR-2025-00058, *Application of Virginia Electric and Power Company for a 2025 biennial review of the rates, terms and conditions for the provision of generation, distribution and transmission services, pursuant to section 56-585.1 A of the Code of Virginia.*

[17] See, *In the matter of the Application of Ohio Power Company for New Tariffs Related to Data Centers and Mobile Data Centers*, Ohio Public Utility Commission ("OPUC"), Case No. 24-508-EL-ATA.

AEP's pipeline of interconnection requests for data centers fell from a projected 30 GW to 13 GW, according to press reports.[18] General concerns have been raised recently in Ohio state regulatory proceedings not yet decided by Ohio state regulators about AEP's over-forecasting of electric load and the seeming lack of sufficient PJM vetting when translated to the wholesale level.[19]

c. Other utility large load adjustments

Like Dominion, other TOs or LSEs reporting significant LLA/data center load growth do not appear to have provisions in place to harness their projections:

- ComEd is currently in a pending proceeding before the Illinois Commerce Commission (ICC) proposing a large load/data center tariff.[20] Its requirements are still not determined. Moreover, there is no disclosure in the matters posted by PJM to support ComEd's LLA regarding whether the terms of this tariff are comparable with those adopted by AEP and Dominion and whether the prospective effect of the tariff will have any impact on ComEd's LLA requests.

- PPL in its filing in support of its LLA request states that it has contractual commitments from its prospective data center customers to support its project load estimates. The filing, however, does not disclose, other than in very general terms, the level of commitment or the maturity of the projects to support PPL's load estimates.

In its review of the LLAs in the current cycle for preparation of the 2026 PJM annual load report, it is incumbent upon PJM to better analyze the contractual arrangements, the level of contractual and financial commitment, and the maturity of each of the data center projects included in the TOs LLAs. An ambiguous statement that a contractual "commitment" exists is not sufficient, as contracts can have few obligations and be drafted to make them easy to avoid. Given the context described above, the current

[18] Z. Skidmore, *AEP Ohio slashes data center pipeline by more than half – report; Utility said that fall is evidence of its data center tariff working,* Data Center Dynamics (Oct. 1, 2025).

[19] See, *Motion to Intervene, Comments and Motion to open an Investigation of AEP's load forecast of the Ohio Manufacturers' Association*, OPUC Case No. 25-735-EL-RDR in *In the Matter of the Application of the Ohio Power Company to Update its Basic Transmission Rider* (dated Sep. 15, 2025). The OPUC dismissed the motion concluding that it was not timely filed. Order dated Sep. 18, 2025.

[20] ICC, Docket ER25-0677, Commonwealth Edison Company, *Proposal for changes to new service requests for large demand project applicants or customers;* ICC, Docket No. 25-0679, Commonwealth Edison Company, *Proposal for revisions to Rider Distributed System Extensions.*

information is bespoke and wholly inadequate to understand on a uniform basis the level of financial commitment undertaken by individual data center customers in forwarding their load estimates to the TOs/LSEs.

The recent experience reported by AEP, moreover, makes clear that there are, or likely will be, important interactive effects between the new generation of large load retail tariffs focusing on data centers and forecasts of future load—with the prospects for substantial reductions in large load requests in the utilities' respective interconnection queues. It is not clear that Dominion, ComEd, or PPL have conducted any such analysis. Adoption of consistent, uniform contractual commitments by data center customers—including the allocation to data centers of the system expansion costs their new loads cause—should be a pre-requisite to inclusion in the load forecast created at the wholesale level of the LLA requests of an individual TO.[21] Transmission Owners without such verification frameworks should have their large load addition requests discounted or excluded from inclusion in PJM's official forecast.

We are mindful that the individual data center customers and the TOs may raise concerns about the confidentiality of individual project information. PJM should explore with the TOs (and through them with the individual data center customers) methods for disclosure of anonymized data to allow for adequate scrutiny. Absent disclosure of this necessary information to support the LLA load request, at a minimum, removal of the request from inclusion in the LLA addition to the load forecast is warranted.

1.4: Probabilistic Forecast Bands

PJM should develop probabilistic forecast bands that explicitly account for uncertainty in large load additions. Each TO's forecast should be assigned to a probability tier based on the maturity and verifiability of the underlying projects. For example:

- High-confidence band: loads with executed retail tariffs or verified financial commitments and interconnection progress;
- Moderate-confidence band: loads supported by officer letters; or
- Low-confidence band: unverified loads.

[21] Letter to FERC Commissioners from the Electricity Customer Alliance (May 30, 2025), p. 2 ("Some best practices [for electric load forecasting for data centers] could be implemented in the near term, such as the requirement that requests for large load additions to forecasts used in wholesale market and planning processes demonstrate commercial viability or be commitment-backed, similar to approaches already taken in some state retail tariffs.") Citing as an example of such state retail tariffs AEP's Indiana affiliate's large load tariff determined in Settlement Agreement in Indiana-Michigan Power, Case No. 46097, Indiana Utility Regulatory Commission (filed Nov. 22, 2024).

Application of these tiers should be tied to the intended response period of PJM's market constructions. For example, the RPM capacity construct should rely primarily or exclusively on the high-confidence band given it three-year manifestation period, while RTEP may incorporate moderate-confidence scenarios due to the extended planning horizon. The calculations that emerge from this banding process should be complemented by additional forecast discounting as discussed above.

The discounting that PJM applies should also account for uncertainty in the realization of announced projects, borrowing at least in part lessons learned from the process applied in ERCOT.[22] In ERCOT's methodology, load projections are routinely adjusted downward to account for uncertainty in the realization of announced projects, the potential for duplicative service requests, and the attrition historically observed between initial load indications and actual interconnections. PJM should adopt a comparable adjustment mechanism by applying empirically based discount factors to account for:

- the absence of a verified large-load retail tariff or equivalent interconnection approval process,
- limited or no financial commitment by the proposed customer,
- lack of demonstrated site control or permitting progress, and
- historical non-realization rates of similar projects within the TO's service territory.

Applied together, probabilistic banding and ERCOT-style discounting would yield a "least regrets" forecast that appropriately weighs uncertainty and protects customers from stranded costs.

Section 2: PJM should align its LAS process with the CIFP process.

Before receipt of the preliminary 2026 LLAs, PJM had initiated the current CIFP process to address the already monumental challenges facing PJM in meeting PJM's reliability requirement based on the 2025 load forecast. The annual PJM 2025 load forecast is driving the CIFP process, without accounting for added impact of the current round of LLAs that will be incorporated into the 2026 annual load forecast. Yet the CIFP promises to change PJM's rule set for delivery years included in and affected by the 2026 forecast. In fact, the CIFP is likely to impact future service levels and the cost of power and, in turn, for the reasons stated above, the load forecast itself. A reasonable forecast thus necessitates consideration of these potential interactive effects between the LLAs and the CIFP.

[22] See, e.g., ERCOT, *Item 8.1, Long-Term Load Forecast Update (2025-2031) and Methodology Changes* (April 2025).

In this context, failure to jointly consider the CIFP reforms and their impact on the 2026 annual forecast—particularly auctions scheduled for June 2026 BRA and beyond—will lead inevitably to wrongly structured capacity market and transmission planning procurements. Avoiding these infirmities should be a paramount objective; PJM should conduct its stakeholder processes and required FERC filings to accomplish this rather than addressing the load forecast in a silo.

Section 3: PJM should exercise due diligence regarding data center industry sectoral risks.

The enormity of the current data-center land rush, the data center industry's sectoral risk and financial basis, and its implications for future electric load growth requires deeper analysis than PJM's usual approach to load forecasting. The reliability and massive cost issues associated with the unprecedented scale and scope of data center growth requires more rigorous analysis at all stages—including preparing, vetting, and approving the load forecasts. Unfortunately, it is not apparent that any rigorous review has been done of the data center sectoral risks in preparing and vetting the load forecasts in the current annual PJM load forecasting planning cycle.

Section 3.1: Capital investment

Neither PJM nor the TOs/LSEs need to be experts in other complex industries. But the vast investment sums that will be supported by existing electric customers that will be necessary to support the huge increases in estimated electric load growth to serve data centers—already reflected in PJM's 2025 forecast and increased again by the TOs/LSEs' large load requests for the 2026 load forecast—makes certain analysis necessary. Such analysis is necessary for PJM and the individual TOs/LSEs to act prudently and to mitigate or avoid incurring massive stranded costs potentially saddled on existing ratepayers.

As described in greater detail in the appendix, important indicators suggest that the data center industry may be in a bubble phase; if that bubble bursts, the huge future electric loads estimated to serve data centers will not occur and a load forecast anchored myopically by recent current trends alone will badly overestimate demand. For that reason, the issue of data center sectoral risk, including the real possibility of a reduction in load growth flowing from those sectoral risks, is relevant to and should be a central element of the conduct of the forecast during this year's planning cycle.

Greater scrutiny is also warranted because of the incentives that may be infecting the land rush of large load interconnection requests.

- On the one hand, developers in the data center sector, increasingly, see access to the power grid, established by earlier in time service requests to the electric utility, as a critical scarce resource, giving value to the developer's development prospect and ability to exclude competitors.

- On the other hand, the transmission-owning utilities increase their profits with expanded transmission capital investment under the current regulatory regime. Forecasted future data center load means more transmission investment, such that TOs have an incentive not to apply rigor to the forecast, particularly where, as is the case in PJM, half of the costs of major regional transmission expansions are allocated on a foot-print wide basis, so that a big portion of the cost is allocated to other TOs, and supplemental transmission projects often receive little (or no) regulatory scrutiny.

Section 3.2: Data Center Supply Chain Constraints

The data center land-rush and the explosion in capital investment already occurring and projected for the future puts huge stress on the supply chains providing the equipment for functional data center development. Limitations in these supply chains can delay or foreclose data center development.

As an example, London Economics International, Inc. recently issued a report analyzing the supply chain for graphic processing units ("GPUs"), a building block for AI enabling data center infrastructure, and concluded there were not enough of these devices under production to meet world-wide demand.[23] Similar constraints, in addition to the siting and grid-access, foregrounded earlier, likely affect cooling trains, transformers, balance of plant items, and telecommunications interconnections.

While normally a peripheral concern at best in electric load forecasting, given the gargantuan size of the build out necessitated by the forecasted increases in data center load embodied by the LLA requests of the TOs, these matters should also be the subject of inquiry. The TOs, subject to anonymizing the data, as described earlier, should request assurances backed by supporting documentation that their data center customers have adequate supply chain support to anchor their claimed load service requests. Failures to do so should result in the discounting or removal of their load requests from the LLA request.

[23] LEI, *Uncertainty and Upward Bias Are Inherent In Data Center Electricity Demand Projections* (July 7, 2025).

3.3 Technological advancements that reduce energy needs.

The physical infrastructure deployed at data centers to permit support of computing and AI applications is under continuous and rapid change due to technological innovation. As described further in the appendix, NVIDIA, the leading supplier of GPUs, is quickening the pacing of introduction of successive generations of more powerful and efficient GPUs. Each new generation of chip technology brings with it changes in supporting technology, such as different methods of cooling and other "balance of plant" equipment required to make the data center functional. These rapid changes all have important implications for energy usage.

Other areas of innovation particularly focused on or affecting energy usage are in rapid development extend to changes in software code, changes in microchip architecture, reconfiguration of the conduct of AI applications, both geographically and as affecting energy usage, reconfiguration of equipment usage to support AI applications and changes in the balance between "edge" and data center hosting of computational activity,[24]

Estimates of the useful life of the most advanced computer processing technologies currently planned for installation in data centers are under five years; yet the electric grid expansion based on the LLA requests which are anchored by a snap-shot of current technology and its energy usage will take much longer than that to accomplish. By the time the grid expansion is complete, the data center technology it was planned to serve likely will have changed with different energy supply requirements. It is incumbent in these circumstances that PJM's load forecasting undertake analyses and incorporate the results of the impacts of these potential technology changes on energy usage.

A central area of further investigation, within scope for PJM's load forecasting efforts in the current exigent circumstances, is also the prospect for achieving dramatic electric energy efficiencies in data center operation, as well as significant improvements in load flexibility.[25] Investigation of both of these parameters as part of PJM's load forecasting efforts is central to anchoring confidence in its forecast.

---

[24] The literature is vast regarding the pacing and trends in innovation in computer processing technology. See, e.g., Stanford University, *The Stanford Energy Technology Review 2025: A Report on Ten Key Technologies and their Policy Implications* (2025), chaps. 1 and 8; Rahman, R. Leading ML Hardware Becomes 40% More Energy Efficient Each Year, Epoch AI (2024); Nadel, Opportunities to Use Energy Efficiency and Demand Flexibility to Reduce Data Center Energy Use and Peak Demand, ACEEE (Oct. 2025).

[25] See, e.g., Nadel, *Opportunities to Use Energy Efficiency and Demand Flexibility to Reduce Data Center Energy Use and Peak Demand*, ACEEE (Oct. 2025); Joint Research Centre (European Commission), *2025 Best Practice Guidelines for the EU Code of Conduct on Data Centre Energy Efficiency* (2025); Cox et al., GridLab and Telos Energy, *Bringing Data Center Flexibility into Resource Adequacy Planning* (Sep. 2025).

*   *   *

In this context, among the relevant documentation and questions that PJM (and at PJM's direction, the TOs) should be asking of the customers seeking to site data centers within each of the TOs' service areas are the following:

1. What is the total capital investment projected to occur over time for the installation and operation of the data center required to support the levels of electric demand incorporated into the customer's service request?
2. What is the source of the funding for the capital investment in response to question 1.
3. Provide documentation to substantiate the responses to 1 and 2. Provide contractual commitments to provide this funding, including credit rating information regarding the identified funding sources.
4. Identify the equipment the developer intends to deploy at the proposed site? What cooling technology does the developer intend to deploy? What computer processing technology does it intend to deploy? Does the developer have purchase orders indicating that it has a contractual right to the equipment and technology intended to be deployed at the proposed data center site? If so, provide copies of these purchase orders.
5. Does the developer have site control? What is its nature (e.g., options vs. purchased land)? Provide the documentation to prove site control.
6. Does the developer have the land use approvals required to construct the data center?
7. Provide the plans identifying the electric draw from the facilities intended to be deployed at the site, certified by a professional engineer.
8. Identify any similar facilities under development within the PJM footprint by the same developer or affiliates and the business plan showing the level of development maturity of these alternate facilities (including site control, financing, and equipment procurement).
9. What is the track record of the developer for data center development? What projects have been completed by the developer? What were their size and attributes (location, electric capacity and demand, and financing requirements)? What was the time-line for completion of these projects? What data center projects are under development by the developer (and their size and attributes)? What data center projects has the developer cancelled or terminated?
10. What measures, consistent with best practices, has the data center developer taken to maximize load flexibility measures?

11. What measures, consistent with best practices, has the data center developer taken to maximize efficiency in usage of electricity of the data center?

As noted above, to protect the confidentiality of the information, PJM (and the TOs) could designate an independent expert third-party clearing house to receive and aggregate the information at a sufficiently granularized level for anonymized public review. Given the systemic implications of the data center load requests, the public interest warrants disclosure of this information.

Conclusion

This year's TOs' LLA requests, primarily if not entirely comprised of future projected data center loads, are gargantuan, building on already massive increases reflected in last year's annual forecasting cycle. They have the attributes of the 1889 Oklahoma land rush, this time propelled by data center developers seeking grid access. Yet the maturity and lack of duplication or speculation embodied in these requests is not in the public record.

This land rush to assert claims to grid access arises, at least in part, from the huge financial bets in the face of gigantic risks placed by and currently driving the data center infrastructure sector in pursuit of its strategic need for scarce resources, including access to the electric grid. The land rush is further enabled by the incentives facing the TOs submitting the LLA requests to build massive transmission facilities, where the risks of electric facility overbuild and stranded costs if the loads do not materialize can be shifted to captive ratepayers.

The PJM load forecast, given its vital role in PJM capacity procurements and transmission expansion planning, must be enhanced—*during the current load forecasting cycle*—to impose greater rigor on the TOs and LSEs to counterbalance the perverse incentives engendering this land-rush. Absent such timely and needed action and immediate reform of its load forecasting process, PJM runs the risk of saddling its existing customers with massive stranded costs in future PJM capacity market procurements and transmission facility expansions.

Appendix – Extraordinary Risks Emanating from the Data Center Sector

The data center "land rush" is catalyzed by and responds to the current trajectory of massive capital investments directed to creating the physical infrastructure in data centers to support the known and anticipated greatly expanded deployment of artificial intelligence (AI) and its seeming voracious need for greater compute capacity.

A critical, but only one among several, supply chain bottlenecks needing resolution to enable further growth at the same or more rapid trajectory for this expansion is the capacity of the existing electric grid to support this expansion and the feasibility of scaling up the grid rapidly in line with the data center land rush's implications for electric supply.

The pinch-point presented by the electric grid's capacity reflects the current trend in data center and computing architecture development for massive increases in the density of electric use (per kW requirements) and very bullish views of the "rebound" effect (or "Jevons Paradox") on the demand for compute.

This is so even if from 2000 to 2018, a period entailing logarithmic increases in compute demand, the always present and then also approaching trade-off (between increased compute and electric usage) was successfully addressed without material increases in electric consumption, due to the always rapid pace of innovation in computing technology.[26] While not dispositive, this prior history of huge technical innovative capacity in the computing industry and the continuing possibilities flowing from this innovative creativity should help inform any current analysis and extrapolations into the future of the electric demands of the data center sector.

There are increasing questions, however, from industry insiders and the financial investment community about the viability of the data center land rush and the continuation of its accelerated growth.[27] Just this month, even the Bank of England in

---

[26] Masanet, E., Shehabi, A., Lei, N., Smith, S., & Koomey, J. (2020). *Recalibrating global data center energy-use estimates.* Science, 367(6481), 984–986.

[27] *See, e.g.,* David Cahn, *AI's $600B Question*, Sequoia Capital (June 20, 2024) (AI annual revenue required to support data center capex and opex at estimated 2024 rates is $600Bn; meanwhile, AI revenue (from OpenAI at $3.4Bn and other AI providers) generating an estimated $100 Bn, leaving a "hole" (needed additional revenues) of $500Bn in annual revenues to justify the data center investment); Dan Robinson, *Alibaba exec warns of overheating AI infrastructure market. Joe Tsai says speculative datacenter builds could exceed actual demand*, The Register (Mar. 25, 2025) (citing Mr. Tsai's comments; also noting that "[i]n January, Gartner said businesses were uncertain about the benefits of investing in AI and CIOs were beginning to question the assumption that AI will lead to a transformation, forecasting that 2025 'is going to be the year of the slide'"; noting also Omdia's estimates for capex spend in 2025 at $576bn, may be lumping capex and opex, inflating the estimates); Claire Burch*, Economics of the AI Build-Out*, Deep Dive, Contrary Research (Aug. 28, 2025) ("AI is

its Financial Policy Committee meeting looking at systemic financial risks raised questions about these developments and their adverse global implications.[28] The major hyperscalers, Amazon, Meta, Alphabet, Microsoft and Google (now sometimes expanded to include Oracle), have expended and, for bullish observers discounting these questions, are projected to expend additional huge sums in fixed investment in data centers over the next five years.[29] As of April, 2025, the biggest four operators control 42% of US data center capacity.[30] Much of that investment entails the acquisition of increasingly powerful computing systems, particularly from NVIDIA, the dominant and leading firm designing and supplying graphic processing units ("GPUs"). State of the art GPUs are the currently foundational equipment enabling the rapid and massive computer processing required to support AI.

---

now the US economy's single largest driver of growth, but it is built on economics that are fundamentally unstable as of 2025. Investors and technologists alike are reconsidering whether AI progress and investment can continue at the same breakneck pace: will efficiency gains, new hardware paradigms, and novel energy sources keep the boom going? Or will the mismatch between falling unit costs and ballooning total costs eventually break the business model entirely?").

[28] *See*, Bank of England, *Record of the Financial Policy Committee on Oct. 2, 2025* (Oct. 8, 2025) ("The risk of a sharp market correction had increased....On a number of measures, equity market valuations appear stretched, particularly for technologies companies focused on Artificial Intelligence (AI).... the Committee noted the future outlook for valuations was uncertain, with both downside and upside risks, Downside factors included disappointing AI capability/adoption progress or increased competition, which could drive a re-evaluation of currently expected future earnings. Material bottlenecks to AI progress—from power, data or commodity supply chains – as well as conceptual breakthroughs which change the anticipated AI infrastructure requirements for the development and utilization of powerful AI models could also harm valuations, including for companies whose revenue expectations are derived from high levels of anticipated AI infrastructure investment."); Makortoff, *Bank of England warns of growing risk that AI bubble could burst* (The Guardian) (Oct. 8, 2025).

[29] The sums are uniformly pharaonic and increasing with newer reporting, but the reporting basis is bespoke. *See, e.g.,* Bain & Co., *Technology Report 2025* (Sep. 2025) ((Figure 2, p. 6, capex spending in AI infrastructure, platform and application levels for Amazon, Alphabet, Microsoft and Meta, 2009 to 2025 est. (2025 - $298Bn)); Brown Advisory, *The Data Center Balancing Act: Powering Sustainable AI Growth* (3Qtr 2025 Update), p. 3, Figure 3 showing quarterly capex for Meta, Google and Microsoft, 2025 to 4thQtr 2026 (est.); T. Morgan*, IDC Makes Ebullient AI Spending Forecast out to 2029*, The NextPlatform (Sep. 8, 2025) (forecasting annual global AI spend rising from $430Bn (2025) to $1.3Trn (2029)); McKinsey Quarterly, *The Cost of compute: A $7 Tillion race to scale data centers* (April 2025) (see exhibit 2, projecting global 3.5-7.3 $ Trillion capex expenditure in data center and related infrastructure 2025-2030 depending on the expansion scenario); T. Morgan, *Datacenter Infrastructure Spending is up and forecasts are even higher*, The NextPlatform (Oct. 7, 2024) (reporting on IDC's estimates that global data center infrastructure rose by 54.7% to $62.3bn in the second quarter, 2024, with total worldwide data center infrastructure spending forecasted to hit $230.5bn for 2024).

[30] Kou and Limandbhratha, *Power for AI: Easier Said than Built*, BloombergNEF (April 15, 2025).

This very high level of investment, however, presumes unprecedented scaling up in future sales revenue from providing AI services in order to justify the investment from consumers and businesses utilizing the equipment, anchored by expectations of rapid wide-spread dissemination and the value added and improvements in productivity anticipated to come from AI innovation and deployment. Major doubts, however, exist about whether this scaling is possible or feasible in the short or intermediate term in line with the trajectory of the companies' investment plans (and its extrapolation into the future).

The initial roll-out in 2022 by OpenAI of ChatGPT, utilizing AI, to great fanfare, with rapid wide-spread deployment, provided the catalyst and anchor for this business thesis. However, the rate of mass deployment of and technical improvement in AI has subsequently slowed, even after incurring substantially increased investment to launch newer versions in the hopes of enabling further deployment. OpenAI's GPT-5 is described as only a modest improvement over previous models.[31] OpenAI has never made a profit and reports are that it likely will achieve $12bn in annual gross revenues this year. Meanwhile, as one of the main vehicles for enabling the rapid scaling theses, OpenAI recently signed a nearly $300 billion five-year purchase agreement with Oracle, starting in 2027, for the supply of the physical infrastructure to support its sales. OpenAI's revenues and sales, however, need exponential future growth from their current levels to justify that level of investment.

[31] The Economist, *NVIDIA's $100 Bn bet on OpenAI raises more questions than it answers. What if OpenAI hits a roadblock* (Sep. 22, 2025) ("Though ChatGPT has more than 700m weekly active users, making it by far the most popular AI application, the response to GPT-5, the research lab's latest family of models, has been underwhelming."); Metz and Weise, *What exactly are AI firms attempting to build?* New York Times (Sep. 29, 2025), p. B6 (ChatGPT paying subscribers only comprise 5% of users); Roge Karma, *Just How Bad Would an AI Bubble Be?* The Atlantic (Sep. 2025) ("[E]vidence is piling up that AI is failing to deliver in the real world…. Research suggests that the companies trying to incorporate AI have seen virtually no impact on their bottom line. And economists looking for evidence of AI-replaced job displacement have mostly come up empty" citing to recent studies by MIT [see below], McKinsey and by Model Evaluation and Threat Research (METR) finding minimal benefits to businesses utilizing AI affecting their profits and for the METR study a decrease, rather than an improvement, in software coding productivity); MIT NANDA, *The GenAI Divide, State of AI in Business 2025* (July 2025) p. 3 ("Despite $30–40 billion in enterprise investment into GenAI, this report uncovers a surprising result in that 95% of organizations are getting zero return…. Tools like ChatGPT and Copilot are widely adopted. Over 80 percent of organizations have explored or piloted them, and nearly 40 percent report deployment. But these tools primarily enhance individual productivity, not P&L performance. Meanwhile, enterprise grade systems, custom or vendor-sold, are being quietly rejected. Sixty percent of organizations evaluated such tools, but only 20 percent reached pilot stage and just 5 percent reached production. Most fail due to brittle workflows, lack of contextual learning, and misalignment with day-to-day operations.").

Increasingly (and particularly over the last year and a half), the sources of funding and risk allocation in the data center industry or "eco-system" have shifted, addressing the uncertainty about achieving the astronomical future scaling of AI applications and compute demand necessary and huge amounts of investment under optimistic scenarios to sustain the business model. The hyperscalers are increasingly turning to more risky debt, rather than cash flow, to fund their capital investment.[32] Recent investment advisory commentary concludes that "hyperscaler free cash-flow growth has turned negative."[33] Moreover, a new family of smaller, highly leveraged companies (the "Neoclouds") have sprung up to provide data center services, both to the hyperscalers, thereby spreading the hyperscalers' risk to the extent they allow the hyperscalers to leverage use of the neoclouds' platforms instead of their own direct investment and balance sheet in participating in data center capacity expansion, and to the broader market of users.[34]

The neoclouds, in turn, are funded by private equity firms, by the hyperscalers, and by entities, particularly NVIDIA, in the up-stream data center supply chain, directly through equity investment and/or enabling collateralization of loans to fund the purchase of the chips supplied by NVIDIA. Questions have been raised about the heightened level of risk arising from this "circular" financing among Silicon Valley entities. In a possible virtuous circle, the financing drives further sales by NVIDIA to enable ultimately supporting the hoped for astronomical growth in scale of AI application usage. On the other hand, it presents also the possibility of a down-ward spiral and collapse in value of the data center investment if the growth in scale in final use does not occur, exacerbated by NVIDIA's rapid pacing of chip innovation, accelerating the obsolescence of its previously sold generations of chips and their value as collateral to the lending.

A recent sign of the increasing risk in the data center eco-system world (sustained for now by continued exuberance in parts of the investment community) and illustrating the circular financing driving these deals is the announced OpenAI-Oracle deal, referenced above entailing a $300 billion 5-year commitment to Oracle from OpenAI, with its current financials wholly disproportionate to the level of investment, and Oracle's ability to treat the investment as remaining permanent obligations (RPOs) in its current

[32] The Economist, *Who will pay for the trillion-dollar AI Boom* (Jul. 31, 2025) (shift in funding sources for data center (AI) investment from hyperscaler cash flow to debt and "a rush to borrow [that is more furious among big tech's challengers [*e.g.,* CoreWeave and other neoclouds]); Gallagher, *Debt is Fueling Next Wave of AI Boom; for aspiring AI players like Oracle much rides on leverage and hope.* WSJ (Sep. 30, 2025), p. B10.

[33] Morgan Stanley, *The GIC Weekly* (Sep. 29, 2025).

[34] Elongated Musk, *Neo-Cloud Economics and Viability in 2025* (Jul. 15, 2025)(description of neo-clouds, share of AI compute capacity supplied by neo-clouds moving from 8% in 2023 to 18% in 2025).

financial reporting. Connecting all the legs of the Oracle-Open AI arrangement requires including the separately announced NVIDIA phased investment of $100 billion in OpenAI, which illustrates the circular nature of what looks like vendor financing. OpenAI buys chips and data center service capacity from Oracle, Oracle purchases the equipment from NVIDIA, and NVIDIA provides the funding to OpenAI to buy the chip services from Oracle.[35] NVIDIA has other circular financing arrangements propping up chip sales to neocloud companies like CoreWeave.[36] These are explained partially as a defensive move to expand sales of GPUs in response to the hyperscalers moves to develop their own GPUs, in turn a response to reduce NVIDIA's pricing power over state of the art GPUs. These arrangements, individually and collectively, puff out dramatically data center expansion and, in front of that the issues confronted by PJM's LSEs, speculative prospecting for sites and scarce grid access for data center development.

Recent business and investment advisory commentary scrutinizes the mis-match between the level of investment in data centers and the level of future revenues required

---

[35] *NVIDIA's $100bn bet on OpenAI raises more questions than it answers. What if OpenAI hits a roadblock?* The Economist (Sep. 22, 2025); Berber and Whelan*, Nvidia Plans Big Open AI Outlay*, WSJ (Sep. 23, 2025), pp. A1 and A2 (NVIDIA to invest up to $100B to support data center and power capacity allowing OpenAI to build and deploy at least 10 GWs of NVIDIA systems for AI data centers. Follows on $300B Oracle-OpenAI deal. Investments will be staged as OpenAI progresses. "That could allow [NVIDIA] to hedge its risk should OpenAI not be able to continue growing at its current rate." "OpenAI will use the cash from NVIDIA's investments to help pay for the chips produced by NVIDIA, a circular arrangement that allows the chip company to turn its balance sheet cash flow into new revenue. Such circular arrangements are common in the AI world and have raised questions about the extent to which new sales reflect genuine market demand versus capital recycled within the industry."). Gallagher, *Debt is Fueling Next Wave of AI Boom; for aspiring AI players like Oracle much rides on leverage and hope.* WSJ (Sep. 30, 2025), p. B10 ("Despite the hype around the company, OpenAI's position looks tenuous. It would need to grow to more than $300 billion of annual revenue in 2030 to justify the spending envisioned in the Oracle contract, DA Davidson analyst Gil Luria estimates – a big rise from the company's current run rate [of annual revenue] of $12 billion."); The Economist, *NVIDIA and OpenAI. Circular Thinking* (Sep 27-Oct 3, 2025), pp. 55-56; Forgash and Ghosh, *Open AI, Nvidia Fuel $1 Trillion AI Market with Web of Circular Deals*, Bloomberg (Oct. 8, 2025).

[36] Rich Duprey, *Are NVIDIA's Latest AI Bids the Biggest Risk to Its $4.3 Trillion Valuation?,* 24/7 WallStreet (Sep. 26, 2025) ("The recent deals between Nvidia, CoreWeave, and OpenAI raise eyebrows due to their interdependent structure. Nvidia's $6.3 billion commitment to buy CoreWeave's unsold cloud capacity through 2032, alongside CoreWeave's multi-billion-dollar deal with OpenAI, create a financial loop where each company supports the others' growth. Some analysts argue this isn't a major issue, suggesting these agreements ensure resource allocation and stabilize revenue amid AI's rapid expansion. However, their circularity invites scrutiny. If Nvidia's purchases are less about market demand and more about propping up CoreWeave's valuation — pegged by some analysts at $75 billion with a $19 billion cash burn rate — it could inflate stock prices artificially, a concern echoed by skeptics who see parallels to past tech bubbles."); *See also*, Kerrisdale Capital Management, *CoreWeave, Inc. (CRWV), Artificial Returns* (Sep. 2025).

to support that investment and is sounding the alarm. Bain and Co. Consulting, in its recent assessment, indicates that meeting the anticipated AI demand by 2030 will require $2 trillion in annual revenue to provide sufficient returns to support the necessary computing infrastructure; "however, even with AI-related savings, the world is still $800 billion short to keep pace with the demand."[37]

This timing mis-match between the massive investment in data centers and the future huge increases in genuine final sales revenues, themselves highly uncertain, is obscured by the hyperscalers' financial reporting practices for depreciation. A major component of the massive AI investment in data centers is the chip and server technology, particularly GPUs, that provide the data center's computing power. Chip and server technology has always had a relatively short useful product life (extending from four to six years), given the continuous rapid innovation of the technology. The pace of innovation in the most advanced chips installed in data centers has accelerated due to NVIDIA's dominance in GPU supply and its recent and announced practice to roll-out a new, much more technically advanced chip every year rather than bi-annually, leaping ahead of the compute power provided by the prior year's chip. This practice by the sector's leading equipment supplier creates further pressure to shorten the useful life of existing legacy equipment.

Given the relatively short useful life of the chip and server technology, magnified by the acutely capital intensive nature of the sector, choice of the depreciation period for the equipment has an outsized impact on the sector as reflected in its financial reporting. Between 2021-2022, Microsoft, Alphabet and Meta, swimming against the current of increasingly rapid innovation embodied in newer investment, extended the useful lives of their server equipment to 4 years. The next year the hyperscalers extended the technology's depreciable lives to 5-6 years, thereby arguably distorting the genuine and economic reality of their investments while unduly enhancing investor optics about the

---

[37] *See,* Bain & Co., *Technology Report* (2025), p. 33 ("Bain's research suggests that building the data centers with the computing power needed to meet that anticipated demand [based on the trajectory of current usage] would require about $500 billion of capital investment each year, a staggering sum that far exceeds any anticipated or imagined government subsidies. This suggests that the private sector would need to generate enough new revenue to fund the power upgrade. How much is that? Bain's analysis of sustainable ratios of capex to revenue for cloud service providers suggests that $500 billion of annual capex corresponds to $2 trillion in annual revenue. What could fund this $2 trillion every year? If companies shifted all of their on-premise IT budgets to cloud and also reinvested the savings anticipated from applying AI in sales, marketing, customer support, and R&D (estimated at about 20% of those budgets) into capital spending on new data centers, the amount would still fall $800 billion short of the revenue needed to fund the full investment."); David Cahn, *AI's $600B Question*, Sequoia Capital (June 20, 2024) (similar question addressing industry context in 2024); Apollo Global Management, *The extreme weight of AI in the S&P 500: Measures of concentration of market cap, returns, earnings and capex* (Sep. 2025).

return on AI initiatives.[38] Shortening of the depreciation period for the technology to three years from the current bespoke reporting of the AI big five, to reflect better the fundamentals of the industry, can have massive impacts on the net worth of the hyperscalers. The Economist estimated this change as knocking $780 billion off the combined value of the AI big five.[39] Similar massive adverse impacts result from such accounting on the value of the chips as collateral to the lending increasingly in use by the data sector to finance its expansion.

In a very recent review of the data center sector and related industries, Morgan Stanley notes the following:

> Capex-cycle maturity can be measured across multiple dimensions, and one of the first we contemplate is inflection of free-cash-flow growth. When cash is abundant and financing costs low, investment can outpace the "technology dream" in a capitalist "arms race." Since 2022, annual hyperscaler capex spending has grown fourfold, depleting cash and halting free-cash-flow growth. In fact, independent research firm Strategas estimates that hyperscaler free cash flow is on track to shrink by about 16% over the next 12 months, based on 2026 forecasts. When free-cash-flow growth for the most richly valued names slows or turns negative, not only do valuations come into question but investor demands for return-on-investment discipline rise.
>
> A second consideration is what's happening among the so-called cash cows—businesses like digital advertising, search, media, streaming and cloud services—that are funding the GenAI build-out. According to a recent report from GQG Partners, over the past four years, most of the hyperscaler cash flow funding GenAI has come from segments with slowing growth due to market saturation, monopolies like search and digital advertising, and areas where competition is increasing.

---

[38] Oscar Mackereth, *Accounting for AI: Financial Accounting Issues and Capital Deployment in the Hyperscaler Landscape* (Cerno Capital Management).

[39] The Economist, *The $4trn accounting puzzle at the heart of the AI cloud. A beancounter's look at the hyperscalers' balance sheets* (Sep. 18, 2025); *See also*, Matthew Fox, *There's a hidden risk lurking in AI stocks in 2025*, Business Insider (Nov. 15, 2024) (similar conclusions citing research from Barclays Bank).

> A final, and most concerning, note pertains to the recent deal-making behavior of GenAI leaders and how it will impact the pace of GenAI build-out. Oracle's strong commitment to build "compute capacity" on behalf of OpenAI will be debt financed and has been pledged to a company with a fledgling revenue model. Nvidia, which already depends excessively on only six hyperscaler customers, has been moving to support Intel foundries and vendor-financing schemes for large language model developers like OpenAI. Like Cisco on one hand and then Tyco and Enron [referring to prior financial "busts" following a boom] on the other, these moves suggest something other than an industry in just the first inning of the build-out.[40]

Here Morgan Stanley signals the bubble nature of the current data center build-out, calling attention to increasing new negative developments affecting the sector—namely "depleting cash and halting free-cash flow growth," "slowing growth due to market saturation," and adoption of increasingly risky financing structures.

Conclusion

The advances to human society potentially enabled by AI and improvements and expansions in compute capacity are enormous. Scaling up of AI usage, based on a snap-shot of current technology's capabilities and, in lock step, commensurate expansion of the electric industry, however, is myopic. Huge and rapid continuing innovations both in AI applications and the physical data center infrastructure are required to support that scaling and the current gaps between projected future and current revenues derived from value added AI usage among the general public.

The electric industry can play a vital role in supporting these developments. However, it should do so in light of: (i) its own very high degree of capital intensity, (ii) much slower ability to expand its physical asset base, (iii) the equitable principles informing its financial support from its existing customer base, and (iv) the very large risks and uncertainties affecting the data center sector.

If, in the course of its rapid, continuing innovations, the data center sector arrives at technical breakthroughs which pivot from the current technology entailing different and lower levels of usage of electricity, the electric industry, and particularly captive utility customers, may suffer. That outcome is far more likely if PJM ties itself to its existing framework for assessing the trajectory of electric usage, which largely accepts projections of massive expansion. The electricity will not be able to pivot with

[40] Morgan Stanley, *The GIC Weekly, It's All About That 'Base'* (Sep. 29, 2025), p. 2.

technological breakthroughs as can the data center section, and will be left with huge amounts of stranded assets.

Similarly, if the huge financial risks currently propelling the data center sector lead to major losses within the sector, captive customers will be at risk of bailing out companies burdened with these stranded assets. Accordingly, the electric industry, as it plans to serve this new kind of customer, should take a measured, and prudent view, meaning more scrutiny at a minimum than accorded presently, of important intermediate and long-term trends propelling the data center sector.