LAS Meeting

August 26th, 2026

Charles River Associates

Meeting agenda

Topic	Timing
Introductions and July 13 LAS Meeting Recap	**9:10 – 9:20 AM**
Top-Down: Unconstrained View	9:20 – 9:35 AM
Top-Down: Constraints	9:35 – 10:00 AM
Bottom-Up Forecast	10:00 – 10:15 AM
Load Profile Assumptions	10:15 – 10:25 AM
Forecast Reconciliation and BTM	10:25 – 10:45 AM
Scenarios	10:45 – 11:00 AM
Break	*11:00– 11:10 AM*
Q&A	11:10 – 11:55 AM
Next Steps & Close	11:55 – 12:00 PM

Today's presenters

Project Director
Margarita Patria, PhD

Project Manager
Oliver Stover, PhD

PJM Data Center Capacity & Peak Demand Forecast

CRA's independent, hourly load forecast for PJM (2027–2047). Presented to the LAS on July 13, 2026

PROJECT SCOPE

Independent PJM-wide data center forecast with hourly load profiles and explicit uncertainty quantification, integrating bottom-up project intelligence with top-down market analysis to support transmission planning, reliability, and resource adequacy.

Horizon	**2027–2047**
Region	**PJM**
Method	**Bottom-up / Top-down**
Timeline	**May–Oct 2026**

NEXT STEPS

Sept 4	LDE form submissions due
Sept–Oct	Individual EDC/LSE meetings
Oct 14	Final results at LAS meeting

METHODOLOGY STEPS

1. **Top-down Demand and Constraint Modeling.** Start with unconstrained demand but forecast what can actually materialize. National AI and non-AI demand is tested against power, semiconductor, infrastructure, and supply-chain constraints. The model is recursive, allowing constraints and market responses to evolve over time but also slowing demand in response to constraints.
2. **Bottom-Up: Site-level project pipeline scored on siting favorability and developer readiness.** CRA reconciles PJM/EDC/LSE and commercial project data, distinguishes true duplicates from development optionality, and evaluates project maturity, ramping, and zonal attractiveness.
3. **Reconciliation:** Top-down determines scale; bottom-up determines location and timing. The two views are reconciled into one internally consistent PJM forecast that honors physical constraints. Excess pipeline may be delayed, relocated, served BTM, ramped more slowly, or not realized; future demand beyond the visible pipeline is represented through generic projects.
4. **Hourly forecasts** reflect evolving data-center operating behavior. Forecasted capacity is translated into hourly load profiles by zone and workload type, incorporating stochastic variation, use of hybrid BTM resources, and changes in load factors and operating patterns over time.
5. **Scenarios** capture alternative paths for data-center growth. Scenarios test different assumptions around AI adoption, efficiency, power availability, infrastructure, and supply-chain constraints to show how demand could evolve under different plausible futures.

Meeting agenda

Topic	Timing
Introductions and July 13 LAS Meeting Recap	9:10 – 9:20 AM
Top-Down: Unconstrained View	**9:20 – 9:35 AM**
Top-Down: Constraints	9:35 – 10:00 AM
Bottom-Up Forecast	10:00 – 10:15 AM
Load Profile Assumptions	10:15 – 10:25 AM
Forecast Reconciliation and BTM	10:25 – 10:45 AM
Scenarios	10:45 – 11:00 AM
Break	*11:00– 11:10 AM*
Q&A	11:10 – 11:55 AM
Next Steps & Close	11:55 – 12:00 PM

No forecasts contained within this document should be considered a "Final" forecast

Top-Down forecast starts with unconstrained national demand, which includes non-AI, AI training, and AI inference components, each with its own growth trajectory

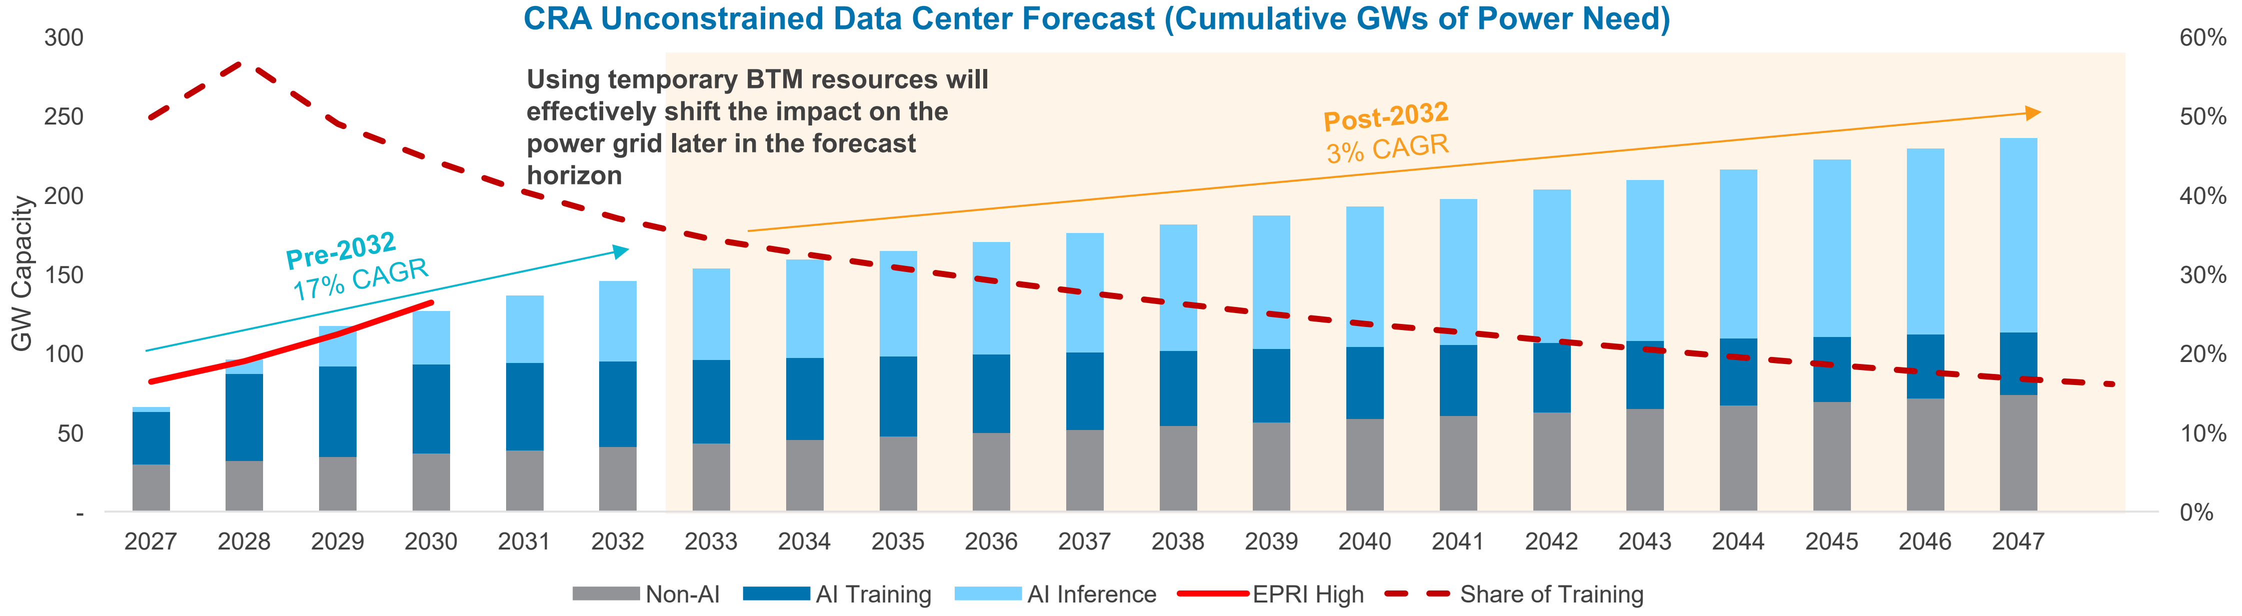

	Pre-2032		Post-2032	
Aggregate CAGR for Unconstrained Demand	**17%**		**3%**	
- **Non-AI Demand**	**+6% YoY**	Assumes linear growth for non-AI related demand.	**+4% YoY**	Assumes linear growth for non-AI related demand.
- **AI Training Demand**	**+10% YoY**	Assumes that public training data is available and expanding compute supply fuel a capability race, driving rapid growth in annual AI training compute demand.	**-2% YoY**	As large models exhaust publicly available data*, a shift to disciplined (smaller) model development and chip processing efficiency drives down total annual compute demand for AI training (see paper)*.
- **AI Inference Demand**	**+75% YoY** *(small base number)*	Assumes that AI adoption grows at a similar trajectory as U.S. internet adoption, with a higher growth factor for both work and individual use. Forecasted edge compute is removed from unconstrained forecast.	**+6% YoY** *(large base number)*	Assumes that AI adoption continues the pre-2032 trajectory. Reaching 90% adoption for individual use around 2032, and 90% adoption for work-related use by around 2041. Forecasted edge compute is removed.

Notes: CAGR = Compound Annual Growth Rate; Sources: Nvidia, Epoch.AI, DatacenterHawk, Pew Research Center, St. Louis Federal Reserve, Ericsson Market Intelligence, Brookings Institute || See Following Slides for Details

Non-AI grows at a steady rate due to traditional digital economy needs

CRA Unconstrained Non-AI Data Center Forecast

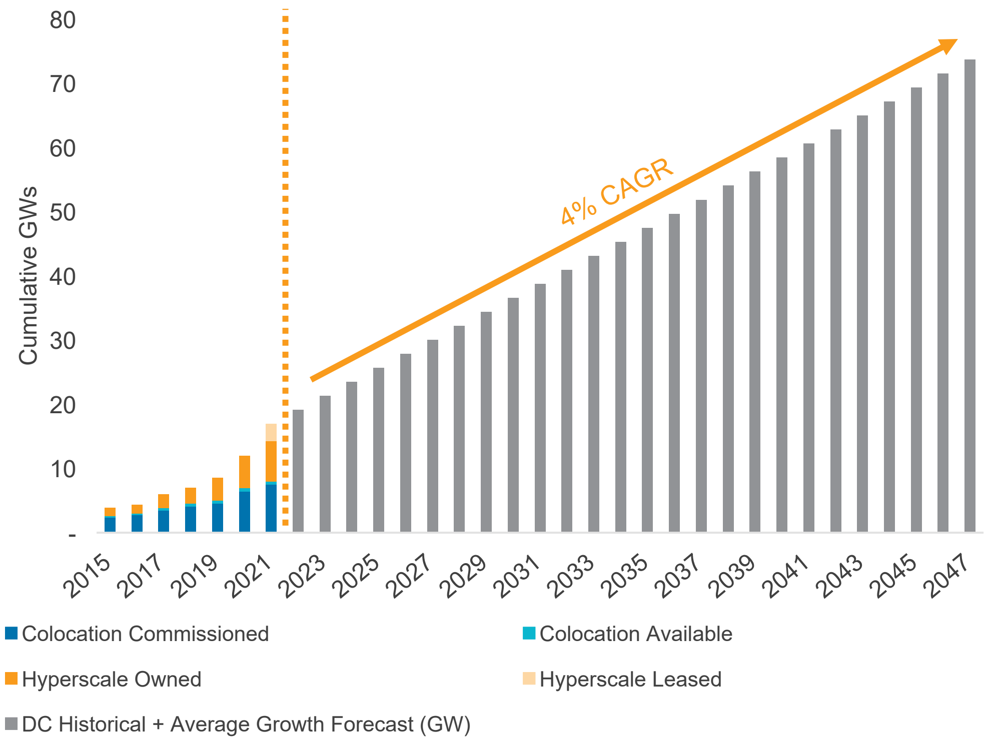

Assumptions

- Assume "AI Growth" occurs in 2022. All-else-equal, assume that non-AI growth stays consistent YoY
- Assumes 'typical' PUE for non-AI data centers to be 1.32 per EPRI (link)
- Elected to use linear model instead of directly applying CAGR growth, as inputting CAGR values may distort long-term growth due to a small starting base

Approach

- Calculate historical GWs growth in Owned, Leased, Available, and Commissioned data centers between 2015 ~ 2021
- Apply average growth GWs to forecast period to capture stable "non-AI" segment of data center growth

Variable	Value	Source
Historical Window	2015 ~ 2021	Investor memos: 2021 cut off based on investment memos of data center providers, mentions "AI" investments first in 2022
YoY Growth	4%	Resulting growth rate from applying *linear* average growth between 2015 ~ 2022

Sources: DatacenterHawk

CRA Charles River Associates

AI training grows, then declines as models exhaust data and gain efficiency

CRA National Unconstrained AI Training Forecast

As Large Language Models (LLMs) exhaust publicly available data, a shift to **disciplined (smaller) model development and chip processing efficiency** drives down total annual compute demand for AI training.

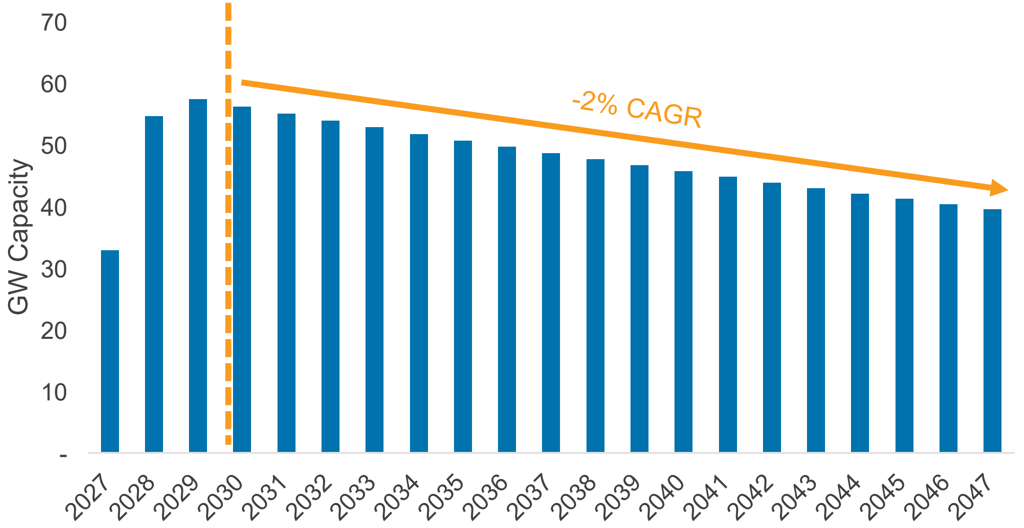

Assumptions

- Assumes fixed server overhead of 1.84 (Epoch.ai), meaning that for every watt of GPU draw, the other components of a rack require 84% of the power of the GPU
- Assumes significant models grow at 11% YoY, using median YoY growth since 1990
- Assumes 1.1 PUE across future periods for AI Training (link)
- Assumes 10% GPU idle rate when computing tasks (link)
- Assumes final training is 10% of R&D, consistent with recent estimates (link)

Approach

- Estimated total training calculations needed using model quantity and average model training intensity
- Converted model training calculations needed to chips required, with assumptions on GPU processing speed and power draw
- Used GPU power draw, supporting server compute needed, and assumed PUE to arrive at capacity needed for new models

Variable	Initial Value	Growth Rate	Source
Avg. Model Size	1.4E+25 FLOPs	**4.3x/year** in recent years, slowing to ½ by 2028, and 1/4th at 7.5% CAGR of growth by 2030 to reflect lack of public data	Epoch.ai: notable models; expected date of public data exhaustion is 2028, with ~95% confidence of exhaustion by 2030
Number of New Models	170 Publicly Notable models	1.11x/year to reflect historical growth rates	Epoch.ai: data on notable models
Chip Processing Speed (FLOP/s)	9.89E+14	1.34x/year growth rate based on average growth of GPUs processing speed*	Epoch.ai: data on machine learning hardware
Chip Power Draw (W)	700 W	1.1x/year growth rate based on average growth of GPUs processing speed	Epoch.ai: data on machine learning hardware

Sources: Epoch.ai, *The Data Center as a Computer (2026 Edition)*, *Small Models are the Future of Agentic AI | Nvidia*,

CRA Charles River Associates

AI inference grows steady as AI is increasingly adopted

CRA National Unconstrained AI Inference Forecast (No Edge)

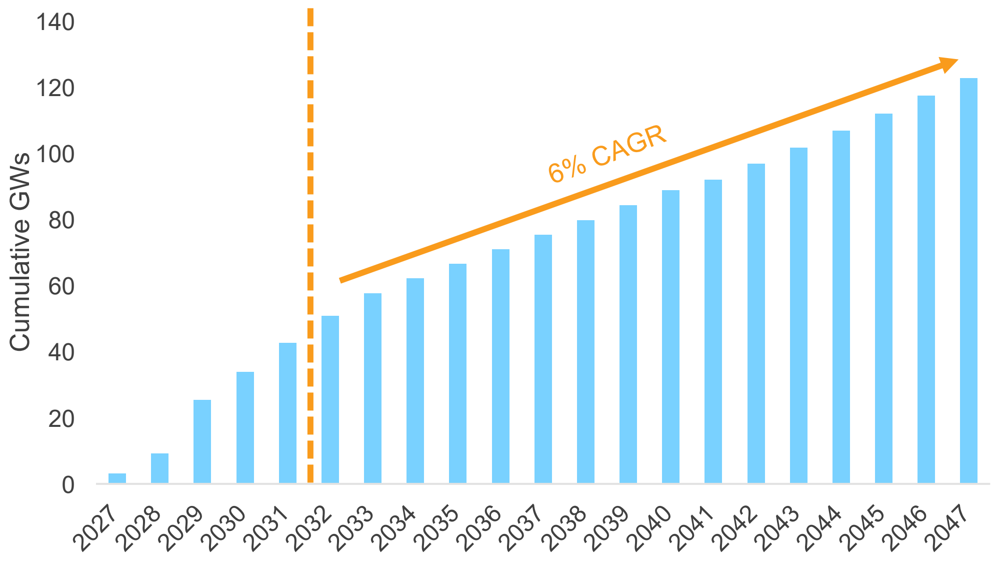

Assumptions

- Assumes edge inference is **not** a part of the large load forecast
- Assumes fixed server overhead of 1.84 (per Epoch.ai), meaning that for every watt of GPU draw, the other components of a rack require 84% of the power of the GPU
- Assumes parameters size scale at the same rate as training size per Chinchilla Scaling Laws and "Small Models" paper, estimated at 7.5% YoY in future years
- Assumes 1.1 PUE across future periods for AI Inference (link)
- Assumes 10% GPU idle rate during tasks computation (link)
- Assumes that "on peak" usage needs to cover 75% of total usage (link)

Approach

- Estimated adoption rates in the U.S. based on historical internet adoption curve and recent work & individual adoption studies
- Combined adoption rates with usage intensity to arrive at total compute needed to process requests
- Converted model training calculations needed to chips required, with assumptions on GPU processing speed and power draw
- Used GPU power draw, supporting server compute needed, and assumed PUE to arrive at capacity needed for new models

Variable	2025 Value	Growth Rate	Sources
Adoption Rates	Around 50% for individual use; 40% for work use	Based on adjusted internet adoption growth rates	St. Louis Fed Study, FRED Internet Adoption Data
Usage Intensity	**19 and 29** from Brookings survey data for individual and work use	Proxy based on mobile internet traffic data growth rates	Brookings Institute AI Use Survey, Ericsson Internet Traffic Growth
Chip Processing Speed (FLOP/s)	9.89E+14	1.34x/year growth rate based on average growth of GPUs processing speed*	Epoch.ai: data on machine learning hardware
Chip Power Draw (W)	700 W	1.1x/year growth rate based on average growth of GPUs processing speed	Epoch.ai: data on machine learning hardware

Sources: Epoch.ai, Brookings Institute, St. Louis Federal Reserve, Ericsson Market Intelligence

CRA anchored AI adoption curve to historical internet adoption shape, while capturing the accelerated rate of growth of AI adoption

FRED Historical Internet Adoption & Modeled Forecast

60% Faster Mid Point Arrival*

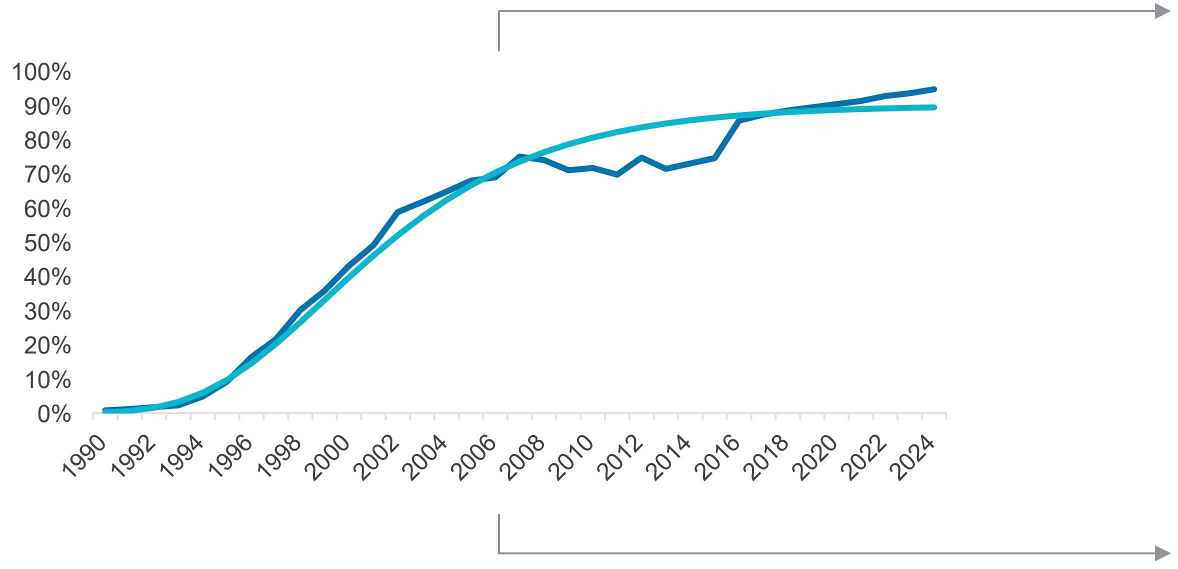

1.5x Growth Parameter, 70% Faster Mid Point Arrival*

*: Start Year adjusted to fit curve, assume AI start year = 2022

Work Related Base Case Adoption Rate Forecast

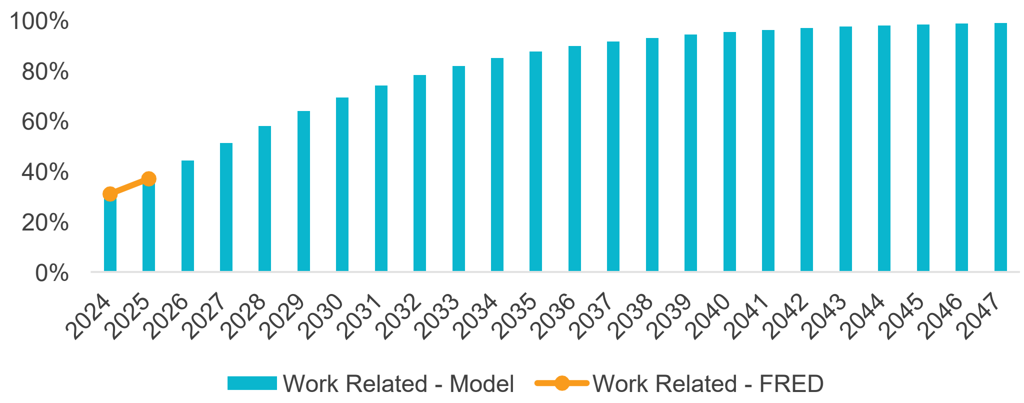

Individual Use Base Case Related Adoption Rate Forecast

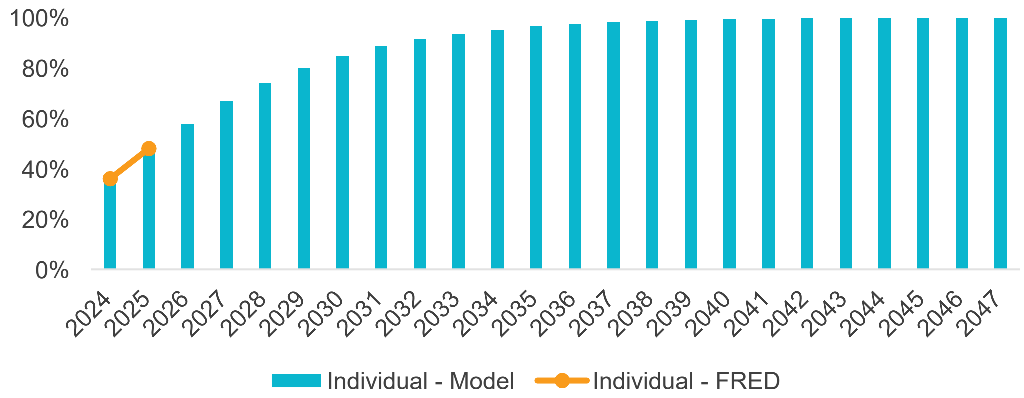

Source: St. Louis Federal Reserve

No forecasts contained within this document should be considered a "Final" forecast

Edge compute growth has been excluded from inference component of the forecast to capture compute migration to local devices

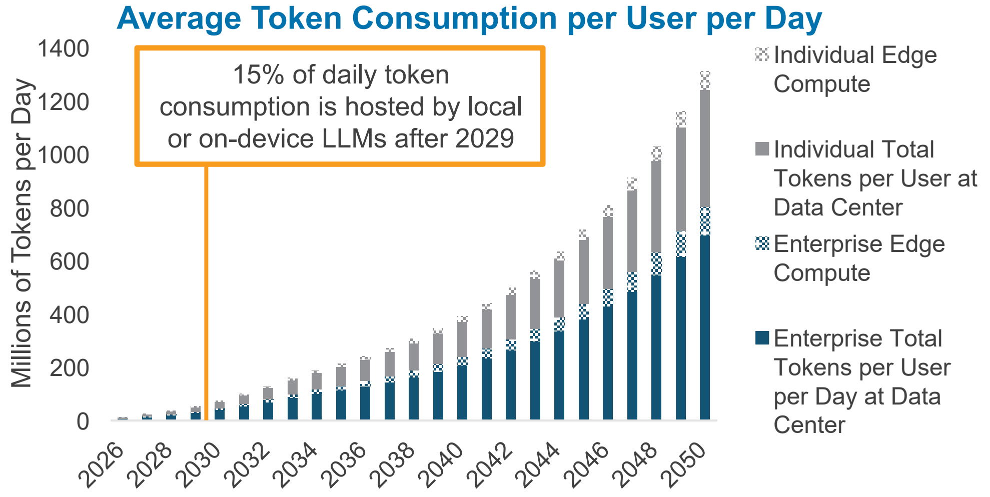

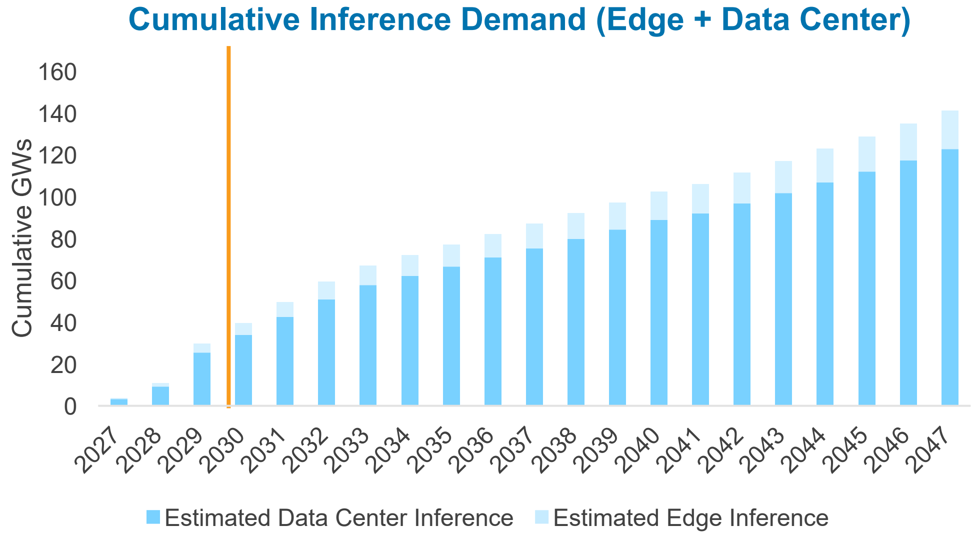

This is because edge compute is projected to shift processing from large data centers into distributed, on-device processing within customer hardware.

Approach

- From the estimated users that have edge-capable devices, remove the share of estimated daily token usage on the edge from total inference demand

Variable	2029 Value	Growth Rate	Source
Share of households with personal computers	90%	N/A	Census.gov data on Computer and Internet Use
Share of tokens moving to local and edge compute	15% of total daily usage by 2029	Assumed constant rate after 2029, linear growth from 2025	OpenRouter State of AI 2025

Assumptions

- Before 2029, local LLM consumption grows linearly to 15% of total average daily consumption
- Assumes that households with edge capable devices will migrate their inference needs onto the edge

Source: US Census Bureau: Computer and Internet Use in the United States: 2021, OpenRouter State of AI 2025 Study, Zima Local LLM Deployment Trend Forecast

CRA Charles River Associates

Historical shares of national AI demand are *initially* held constant across key geographies

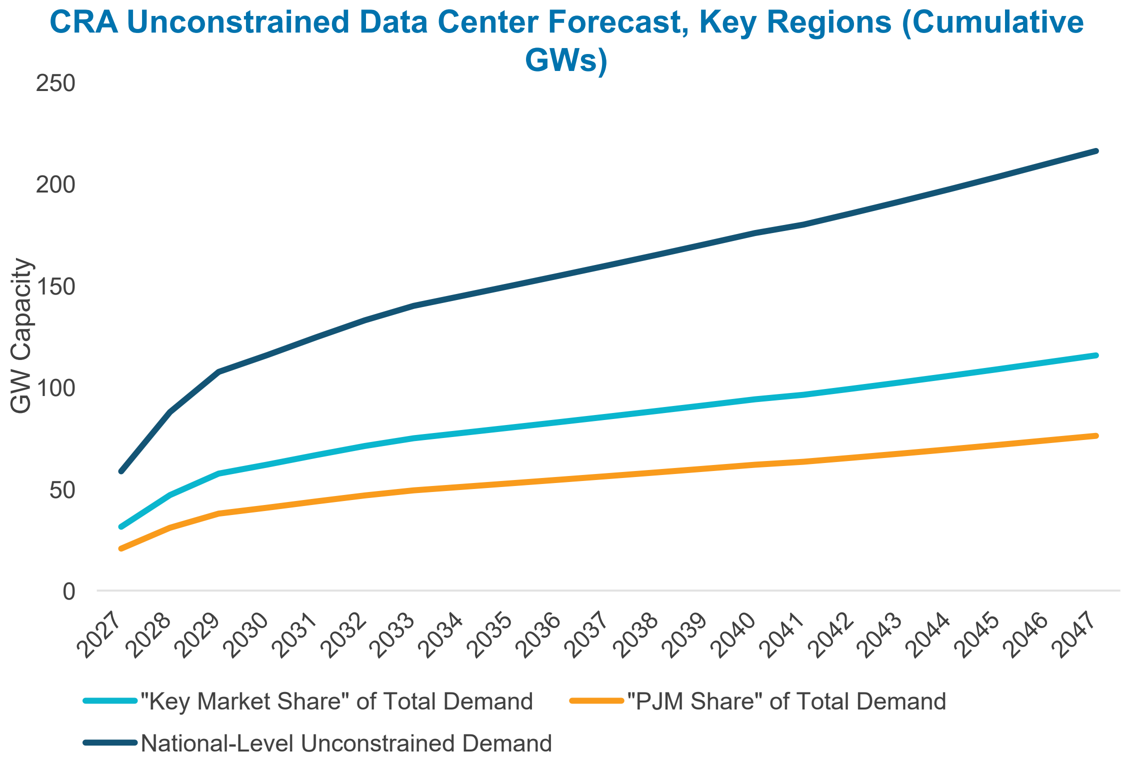

Assumptions

- Assumes market share stays constant across forecast horizon
- "Market Share" is defined as the share of Owned, Operated, and Under Construction GWs in a specific ISO/RTO *(e.g., PJM, MISO, ERCOT)* relative to the total GWs in North America

Approach

- Apply historical market share percentiles to national-level, unconstrained data center forecast
- Assume US data centers do not go international, and PJM's market share remains constant, unless PJM-specific constraints slow growth. Unmet PJM demand may try to shift to other Key Markets listed below.

Variable	Value	Source
"Key Markets" Market Share *(PJM, MISO, ERCOT, SPP, SERC)*	54% (of total market share)	CRA Data Center Database (DatacenterHawk, Cleanview)
PJM Market Share	35% (of total market share)	CRA Data Center Database (DatacenterHawk, Cleanview)

Sources: CRA Data Center Database (DatacenterHawk, Cleanview)

CRA Charles River Associates

Binding constraints will limit the level of realizable demand. As constraints relax, growth rates preserved from the unconstrained forecast will be applied to the constrained forecast

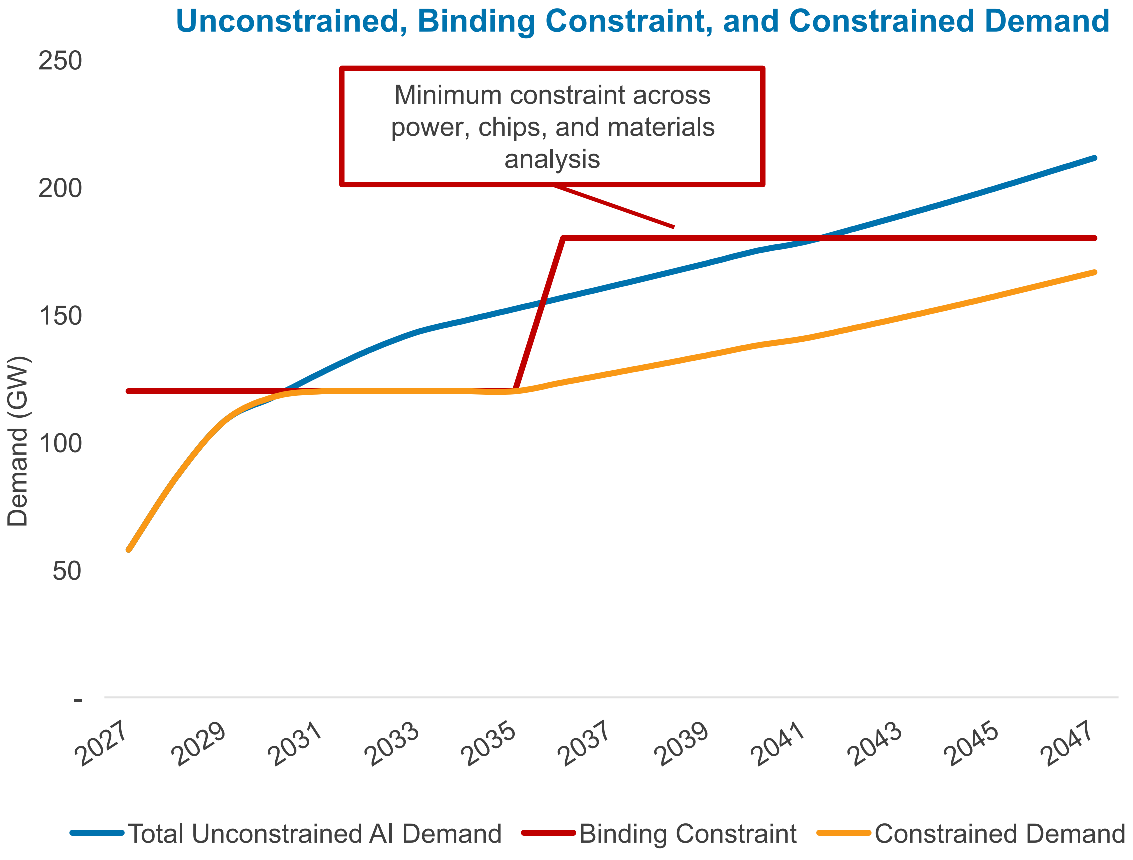

*Binding constraints presented here are illustrative for discussion on constraint logic

Approach

- Apply constraints on unconstrained demand. The constraint acts as a ceiling for how much demand can be realized
- Constraint is minimum across all constraints considered: power, chips, and material availability
- Growth rates (g_t) are preserved from the unconstrained demand, but applied to the lower of prior forecast or realized demand:

$$Target_t = \min(Achieved_{t-1}, Target_{t-1}) * (1 + g_t)$$

Assumptions

- Capacity that could not be built is assumed to be partially lost through efficiency improvements, adoption delays, or price friction, rather than returning as pent-up demand.
- Overshoots do not increase future targets. Functionally, the forecast can be adjusted downward, but not upward.

Meeting agenda

Topic	Timing
Introductions and July 13 LAS Meeting Recap	9:10 – 9:20 AM
Top-Down: Unconstrained View	9:20 – 9:35 AM
Top-Down: Constraints	**9:35 – 10:00 AM**
Bottom-Up Forecast	10:00 – 10:15 AM
Load Profile Assumptions	10:15 – 10:25 AM
Forecast Reconciliation and BTM	10:25 – 10:45 AM
Scenarios	10:45 – 11:00 AM
Break	*11:00– 11:10 AM*
Q&A	11:10 – 11:55 AM
Next Steps & Close	11:55 – 12:00 PM

Few physical constraints are truly PJM-specific; most reflect broader national or global limitations that manifest differently within PJM

Constraint	Geographic Scope of Constraint			Rationale
	PJM	National	Global	
Semiconductor availability	✓	✓	✓	Global manufacturing capacity and supply chains
Critical metals / materials – copper, steel, aluminum	✓	✓	✓	Globally traded commodities and manufacturing inputs
Power availability / ability to serve new load	✓	✓		Increasingly a U.S.-wide constraint, though severity varies by market
Generation / supply build limits – gas turbines, batteries, nuclear, etc.	✓	✓		U.S. deployment constrained by manufacturing capacity, with globally sourced equipment/materials
Generator interconnection queues	✓	✓		National challenge, with PJM-specific queue rules, backlog and timelines
Generation retirements	✓	✓		Occurring across the U.S.; PJM's specific retirement trajectory drives regional impacts
Transmission availability / buildout	✓	✓		Broad U.S. constraint, with PJM-specific transmission needs and planning processes
Natural gas infrastructure / fuel availability	✓	✓		National/regional pipeline constraint, with specific limitations within parts of PJM
Permitting / siting	✓	✓		Federal, state and local requirements constrain development across the U.S., with PJM-state differences
Capacity market / market design	✓	✓		Specific to PJM's market structure and rules

Omdia assessed semiconductor constraint using the following methodology

Integrated analysis combines Omdia's proprietary semiconductor market intelligence with primary and secondary industry research and analyst judgment to assess the extent to which semiconductor availability may constrain U.S. data center development.

Validation: Key assumptions and conclusions are cross-checked against multiple independent sources where possible. Because semiconductor capacity, product mix, yields, customer allocations, and data-center configurations evolve over time and are not always publicly disclosed, results should be interpreted as analytical estimates rather than precise measures of physical capacity.

Modeling: Inputs are synthesized to estimate semiconductor requirements for projected data center deployment and compare against the industry's ability to supply the necessary technologies. Where directly reported data are unavailable, Omdia applies assumptions, industry benchmarks, technology relationships, and expertise to develop reasonable estimates.

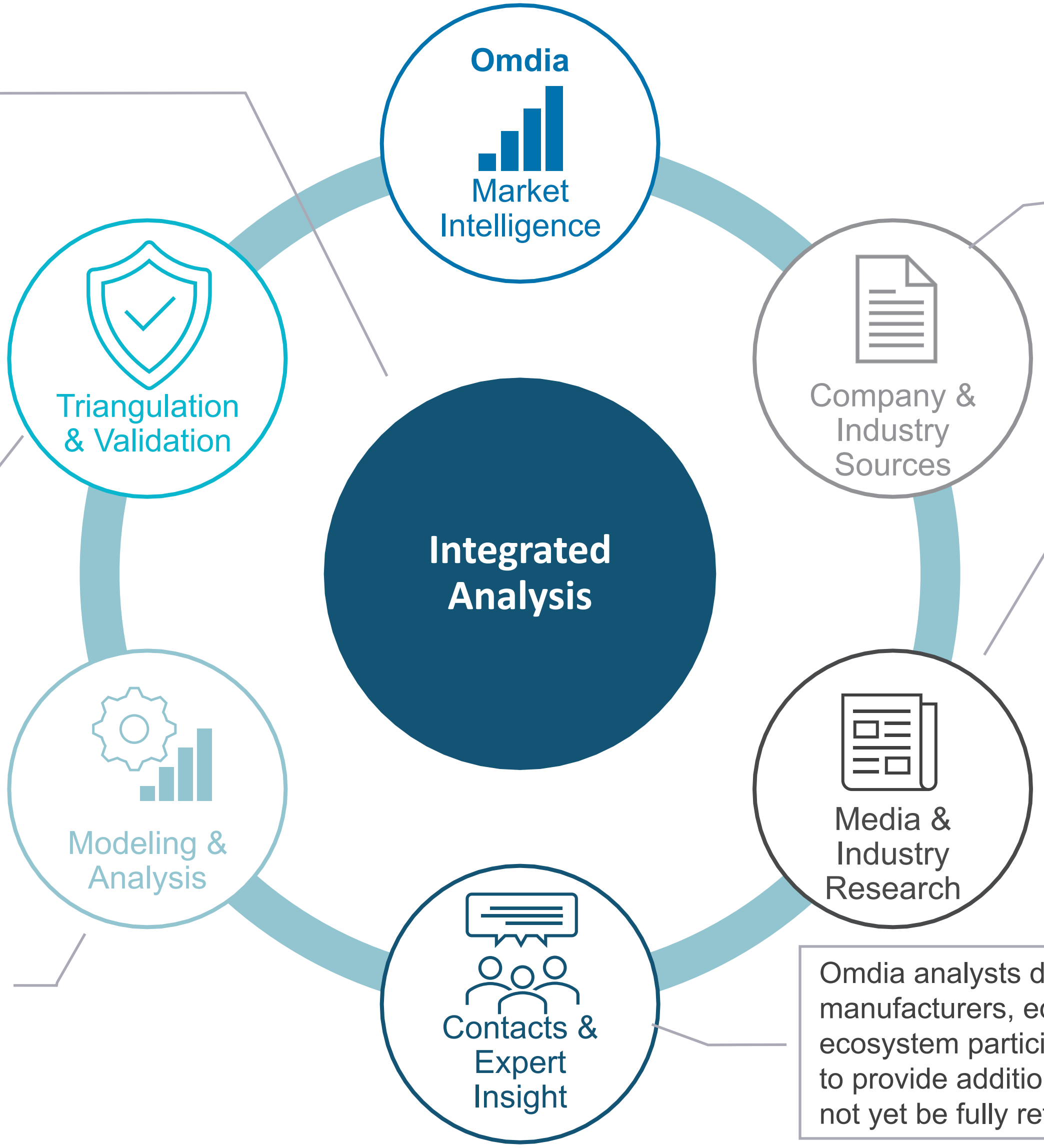

Public company disclosures, earnings commentary, investor presentations, technical documentation, supplier announcements, industry association data, and other published industry research are used to evaluate semiconductor production capacity, technology roadmaps, capacity expansions, and potential supply constraints.

Trade publications, financial media, and other credible reporting are reviewed for information on capacity additions, shortages, lead times, supply-chain developments, customer commitments, and changes in semiconductor and data-center investment plans.

Omdia analysts draw on discussions with semiconductor manufacturers, equipment and materials suppliers, data-center ecosystem participants, customers, and other industry contacts to provide additional context and identify developments that may not yet be fully reflected in published data.

CRA Charles River Associates

Semiconductor constraints are adjusted for the immediate- to near-term and beyond as risks change and bottlenecks emerge and resolve

Term	Critical constraints	How binding?
Immediate-Term (2026-27)	Advanced chip packaging capacity remains constrained	**Moderately** – ramp trails demand, timeline uncertain
	Critical chip packaging material in short supply, significantly increasing costs	**Very** – gap hits 40% by 2028
	Lead times for key semiconductor inputs remain elevated	**Not Very** – new capacity is expected by 3Q27
	High-performance AI memory is fully allocated; limiting new projects	**Moderately** – stack yield limits near-term scaling
Near-Term (2028-32)	Manufacturing is concentrated in a small number of countries	**Moderately** – concentration and geopolitical risk, but does not constrain output on its own
	Limited specialized production environment and workforce	**Partly** – output is limited by facility growth and available workforce, construction and upskilling can take years
	Next gen memory is still uncertain and faces scaling risks	**Very** – manufacturing capacity limits efficiency growth
	Advanced manufacturing technology is not at full-scale	**Not very** – future scaling may be impacted, but long lead time to catch up
	New capabilities are outpacing manufacturing capability and capacity	**Moderately** – manufacturing limits outputs until scaled up which can take years

Long-term and secondary impacts

- Physical and disaster risk in concentrated manufacturing regions
- AI hardware consuming memory supply that would support traditional server deployments
- Emerging difficulties sourcing advanced materials at scale
- Limited testing capacity for advanced AI chips
- Assembly and delivery of complete AI systems is constrained by manufacturing capacity

CRA Charles River Associates

CRU assessed metal constraint using the following methodology

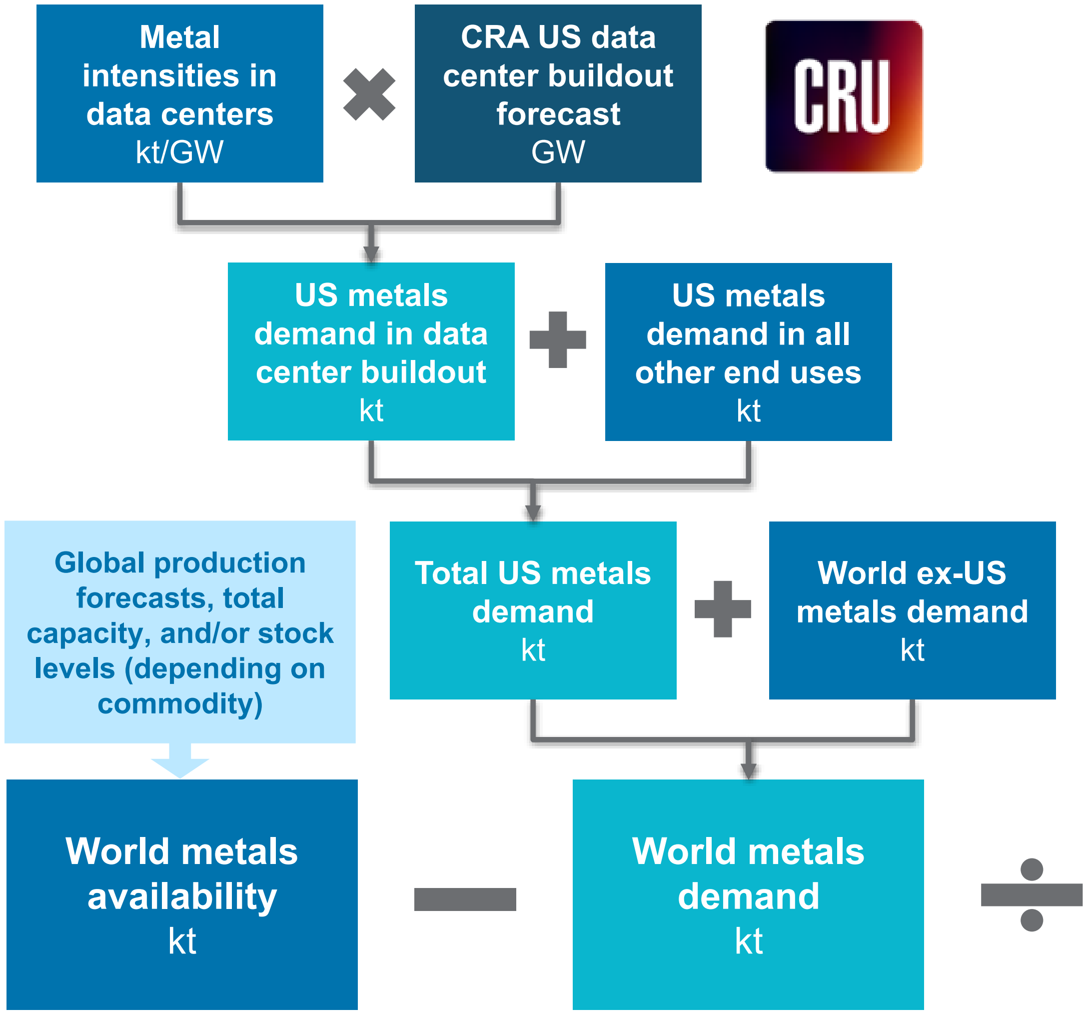

CRU's assessment evaluates whether global production, inventories, and available manufacturing capacity for copper, aluminum, and electrical steel can meet anticipated demand growth from data centers and other end uses, focusing on 2027-2030 as the key years for potential shortages.

CRA incorporates CRU's findings into the US data center forecast by translating any projected global supply deficit into a potential reduction in US buildout. The constraint is allocated based on the US share of projected demand for the affected material, and the resulting material shortfall is converted into an equivalent reduction in data center capacity (GW).

Despite potential global constraints, data centers may be better positioned than many traditional end uses to secure scarce materials, especially given the relatively small portion of total development costs. As a result, actual impacts could be lower than modeled.

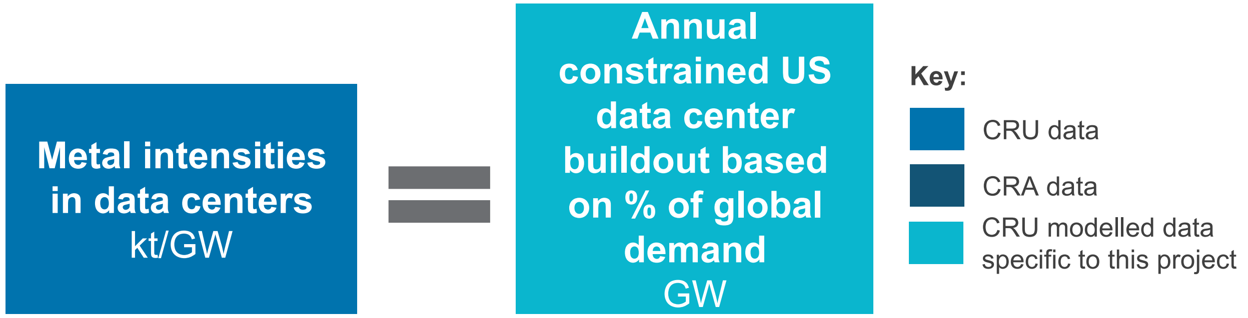

Critical metals and other materials face most risk in the near term, but broader long-term constraints are not expected

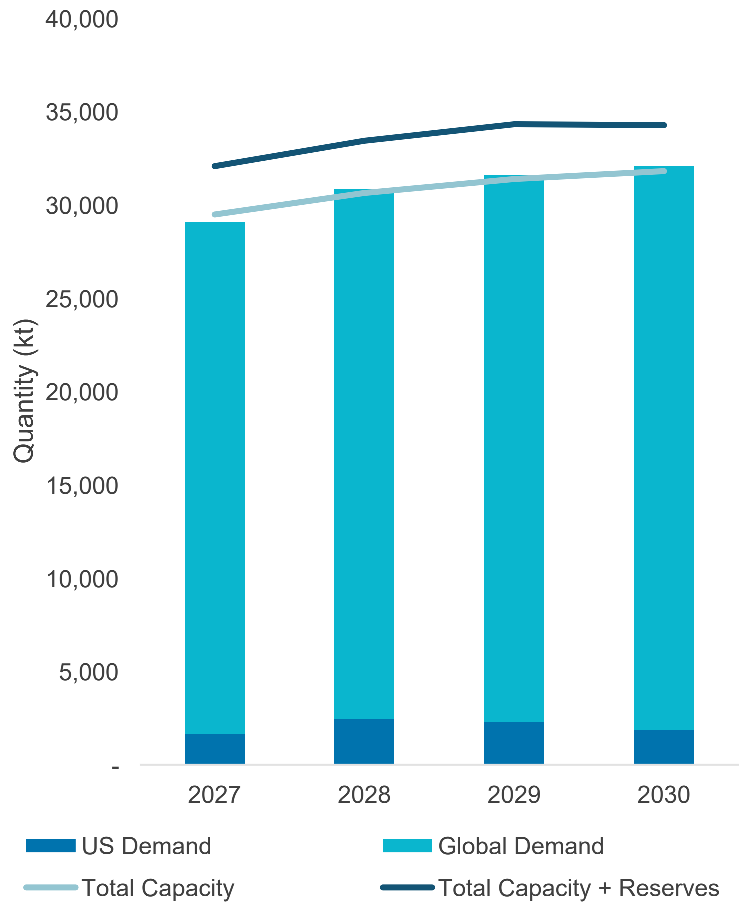

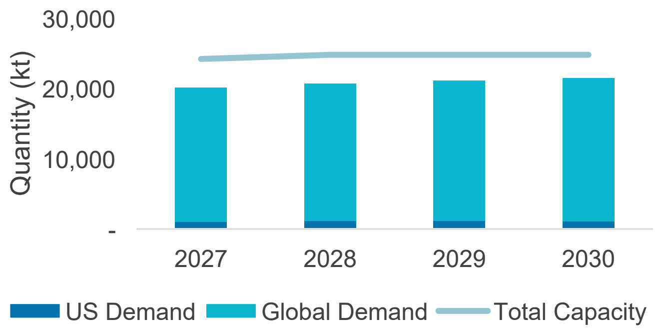

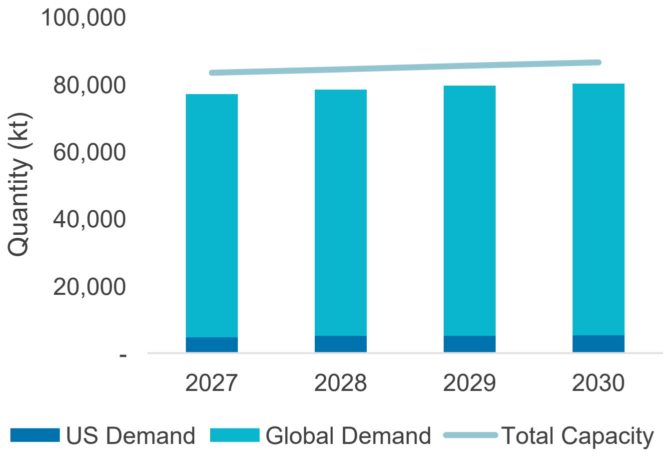

CRU's analysis indicates material availability for data center development should generally be sufficient over the long term, with the primary risk period occurring during the expected peak of global data center buildout from 2027-2030.

Copper faces the greatest potential supply risk, as availability is ultimately constrained by mining production and the pace at which new supply can be developed. For aluminum and electrical steel, the market typically operates with some unused capacity, allowing producers to increase output in response to higher demand. CRU therefore evaluates both current production and capacity that could be brought online to determine whether shortages are likely to emerge.

Beyond 2030, CRU expects global markets to adjust to growing demand through capacity expansions, investment in new production, and other supply-side responses. While short-term imbalances may occur during the peak buildout period, the outlook suggests that material markets can adapt over time, reducing sustained shortage risk.

CRA Charles River Associates

Constraints on power supply effectively slow data center adoption

Reliability Assumption
Reliability standards will be maintained through policy, market, or operational interventions

Near-Term Flexibility
BTM resources (RICE, fuel cells, CTs, CCs, storage, and DR) bridge the gap

Long-Term Evolution
Workloads shifting across facilities and regions are expected to increase over time, improving system flexibility

- CRA assumes that policy, market, or operational actions will return the market to the traditional "**1-in-10**" reliability standard in the longer term. Potential mechanisms to ensure long-term reliability outlooks include state policy interventions, or data center commitments to provide flexibility. Current projected shortfalls provide an upper bound for the tolerance for capacity shortfalls.
- As a result, future data center growth will require increasing levels of load flexibility and/or delayed ramp rates as a result of PJM-wide and zonal supply-demand tightening. While data centers may gain greater ability to shift computing workloads over time, **BTM resources will play a critical near- and medium-term role.**
- Potential BTM resources include **RICE units, fuel cells, combustion turbines, combined-cycle units, batteries, and demand response** from other customers.
- Like front-of-meter resources, BTM resources face limits on **deployment at the site, zonal, and national levels**. However, BTM resources may be deployed more rapidly because they can often avoid traditional generator interconnection processes.
- The ability to **dynamically shift computing demand across sites and regions** is expected to increase over time as the workload shifting technology becomes more mature, enhancing overall system flexibility.

CRA constructed a forecasted view of PJM's electric supply for each year of the study period by PJM zone and technology in unforced capacity ("UCAP") terms

Supply Inputs

Existing Fleet

Begin by taking PJM's existing generation fleet and converting it into UCAP

PJM's existing fleet anchors future supply position while accounting for known retirement dates

Near-Term: PJM Interconnection Queue

Layer in known projects from PJM's queue as the source of new supply

Apply drop-out rate – Based on historical completion rates for the legacy queue and observed Transition Cycle attrition for TC1/TC2 projects (anchored to BRA)

Queue used here, not the model, to avoid letting supply automatically rise to meet demand (circular logic)

Long-Term: Capacity Supply Forecast

CRA used a capacity expansion model to forecast generic supply additions by zone and capture the full extent of long-run resource additions

Model Constraints

Build limits: physical delivery constraints for gas and storage each year

Transmission limits: zonal delivery capacity (CETL vs. CETO)

Output

Year-by-Year, Zone-by-Zone Supply

A zonal UCAP supply forecast used as a constraint on the demand model

Where supply cannot meet load in a given year, the unmet demand is delayed, met behind-the-meter, or relocated to other regions

For illustration, supply is shown against PJM's 2026 load forecast in the following slide

The supply forecast is one of the main constraints applied to the data center demand forecast—it determines how much new load PJM can physically serve, when, and where

Base-case supply forecast (UCAP) vs. PJM 2026 load forecast

There is not enough supply to meet data center growth charted in PJM's **2026 load forecast**; 2030-2035 represents the deepest shortfall

PJM Base Case Supply Forecast (UCAP)

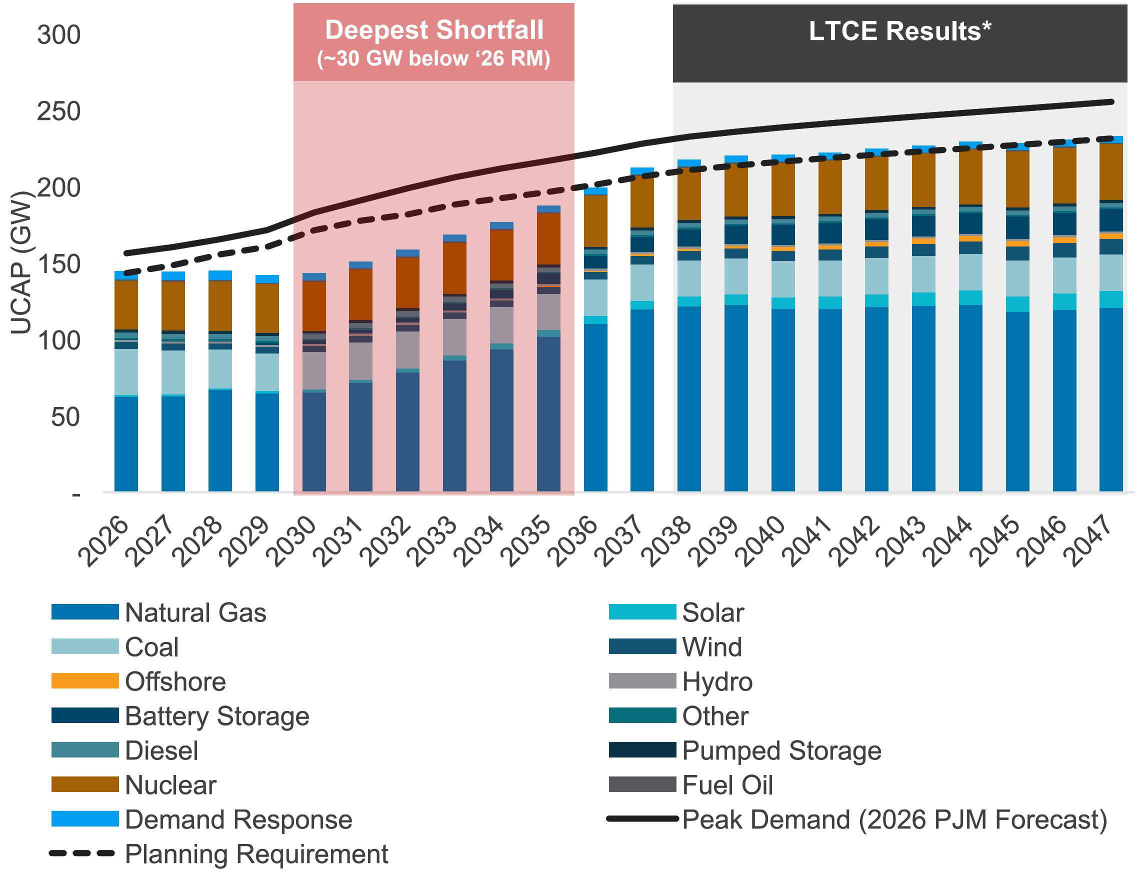

**Note – This view reflects only the LTCE buildout post-2038. It does not incorporate the post-2038 additional gas/storage that can be built up to their annual build limits.*

Key Takeaways:

- **Short from the start:** A structural supply gap opens in 2027/28 and persists across the horizon.
  - Deliverable UCAP falls below the reserve requirement even before major load growth arrives
- **Physical, not financial:** The gap holds despite strong willingness to pay because turbines, transformers, EPC labor, and interconnection throughput cannot be accelerated by price signals alone.
- **Deepest gap in the 2030s:** ~30 GW below the reserve requirement (with 2026 load forecast), when the interconnection queue thins and LTCE builds are subject to the gas and storage build limits. This highlights the impact of global supply chain limitations that PJM cannot control.
- **Load goes elsewhere:** Gap pushes new large loads toward BTM self-supply, relocation to headroom regions like MISO and SPP, waiting for interconnection, or becoming subject to IRAS.
- **Moving target:** The PJM 2026 load forecast shown here is a placeholder; our demand forecast is iterative, so as load and queue status update, the size and timing of the gap will shift rather than remain fixed.

CRA Charles River Associates

Transmission limits (CETO) scale based on known 2030 values

Assumptions

- **Pre-2031 constraints:** CRA assumes published CETO/CETL values are applicable to the beginning of the forecast period.
- **Future CETO (post-2030):** Updated each year by zonal peak-load growth, then adjusted for net in-zone UCAP changes by accounting for capacity additions and retirements.
- **Data center headroom:** The additional load a zone can absorb before its Tx system can no longer reliably deliver power. As load grows, the zone's import need (CETO) rises; once it approaches the zone's import capability (CETL), the zone is constrained and cannot host more load until transmission is upgraded.
  - CETO * 1.15 > CETL is the PJM published constraint formula
- **Interim relief:** CRA does not add additional power to a constrained zone to relieve the deliverability gap before the line is built.
- **Simplifications:** modeled at zonal LDA level only.
  - Sub-zonal LDAs (DPL South, CLEV, PSEG North) captured within parent zone; RECO inherits EMAAC

Post-2038: Transmission constraint no longer applied

CETO after 2030 is projected based on the 2030 value

2030 CETO is scaled by the change in <u>peak coincident load</u> relative to 2030 levels

CETO is then adjusted to reflect changes in future-year and 2030 available capacity

$$CETO_{t,z} = \left[CETO_{2030,z} \times \frac{LOAD_{t,z}}{LOAD_{2030,z}}\right] - (UCAP_{t,z} - UCAP_{2030,z})$$

** Subscript t denotes year of analysis; z represents zone*

Sources	Additional Information
PJM Planning Parameters	For published auction years
2030 RTEP Values	CETO/CETL provided by PJM for upcoming 2030 RTEP

Transmission limits (CETL) and transmission expansion (RTEP) account for the ability to build new transmission

If CETL exceeds CETO x 1.15 in year *t*, the zone is considered **unconstrained**

$$CONSTRAINED_{t,z} = \begin{cases} CETL_{t,z} \geq CETO_{t,z} \times 1.15 = 0 \\ CETL_{t,z} < CETO_{t,z} \times 1.15 = 1 \end{cases}$$

If CETO x 1.15 exceeds the CETL for that year, the zone becomes **constrained**

Feedback welcome!
Results are subject to assumptions around how Tx constraints relax over time (which is in PJM's control). Failing to account for the ability to invest in new Tx would artificially limit growth.

Transmission capacity is assumed to be added through the RTEP process after 3 years of construction, available on year 4

$$CETL_{t,z} = CETL_{t-1,z} + ADDER \times CONSTRAINED_{t-3,z}$$

Calibrated to 4,000 MW based on historical transmission additions

Assumptions

- **RTEP:** If a zone ever becomes constrained, it ignites the RTEP process to bring new transmission online.
- **Future CETL:** Updated each year by zone with the previous year's CETL and any additional transmission.
- **Adder value:** Parameterized as 4,000 MW based on historically observed transmission additions; all additions are assumed at this value.
- **Timeline:** We assume a 4-year timeline from the zonal constraint becoming active to when relief occurs. This assumes that a constraint can be identified a year early with perfect foresight (5 years total) and is standardized across all zones.
- **Compounding additions:** If a zone adds transmission but still is constrained, then the process is assumed to ignite again. The zone can continue adding transmission until after 2038 when the transmission constraint is no longer applied.

Sources	Additional Information
PJM Planning Parameters	For published auction years
2030 RTEP Values	CETO/CETL provided by PJM for upcoming 2030 RTEP

Natural gas build limits (historical) show uptick in demand for turbines in US

Historical Share of Gas Turbine Orders (2001-2025)

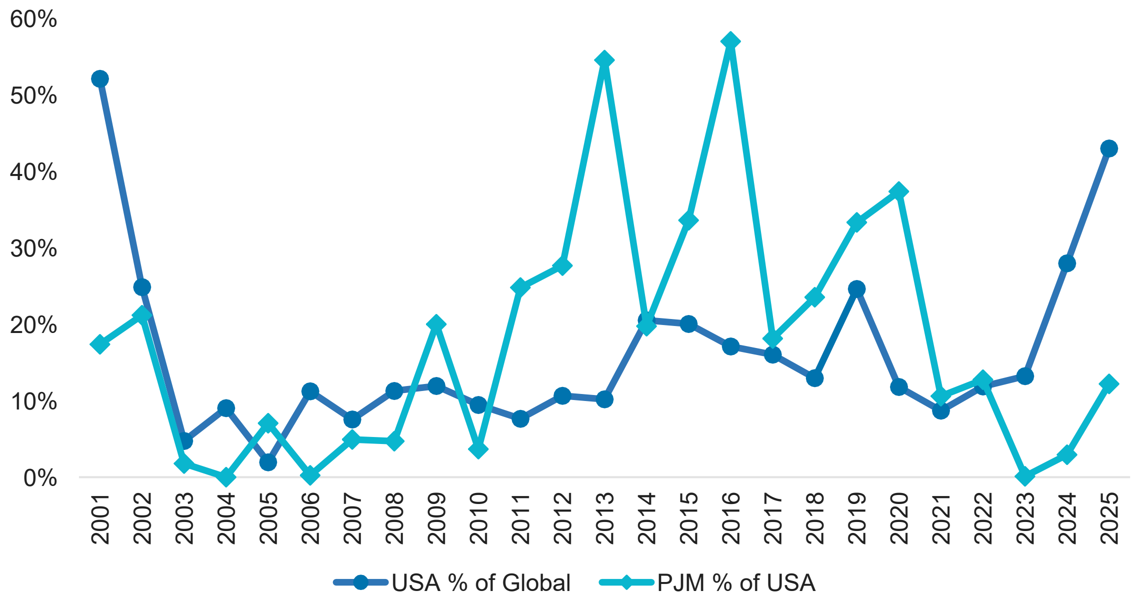

Sources	Additional Information
McCoy Power Reports – Gas Turbine Order Data	Global gas turbine order database (1980-2025), including project location, order and operating dates, turbine model, manufacturer, purchaser, and delivery status.

Assumptions – Natural Gas Build Limits

- Based on 2001-2025 historical gas turbine order and delivery data.
- U.S. share of global orders and PJM share of U.S. orders are held at 2025 levels to reflect sustained data center-driven demand growth.
- Lead-time analysis shows recent gas turbine projects require ~36 months from order to operation, with orders translated into future online capacity accordingly (e.g., 2025 orders → 2028 operations).
- Future limits scale with turbine manufacturing expansion, OEM backlog growth, and delivery trends.
- Build limits represent a physical upper bound based on turbine availability, supply chain constraints, lead times, and EPC execution capacity.

Gas Turbine Deliveries – Historical Trend

- U.S. share of global gas turbine orders increased to ~42% in 2025, reflecting growing domestic power demand.
- PJM represented ~12% of U.S. turbine orders in 2025, with recent growth supported by accelerating data center development.
- Recent ordering trends indicate increasing U.S. and PJM demand for new gas generation, supporting higher future build capability.

CRA Charles River Associates

Natural gas build limits (2027-2047) are assumed to grow over time

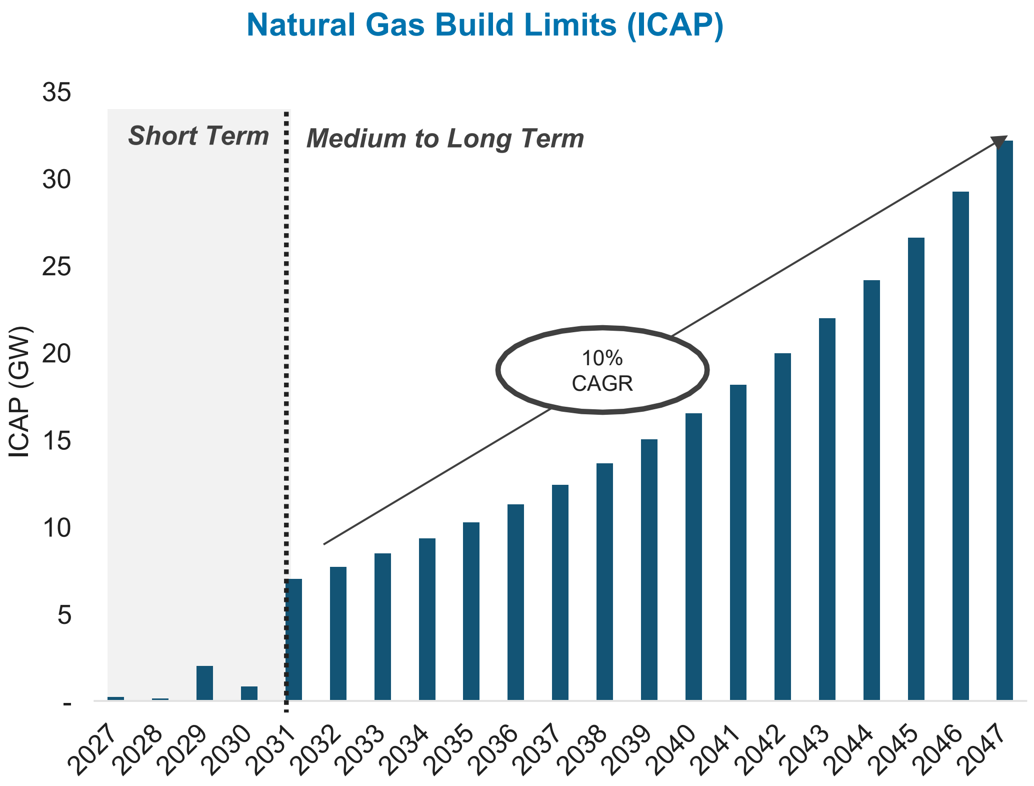

Methodology

Short Term Build Limits (2027-2030)

- Build capability based on active PJM queue projects.
  - Gas projects from queue only included if either under construction or as mentioned in an IRP

Medium to Long Term Build Limits (2031-2047)

- Annual build limits derived from historical gas turbine order trends and projected manufacturing capacity growth.
  - Global gas turbine manufacturing capacity assumed to increase 10% annually based on OEM expansion plans and announced investments*.
    Note: The 10% growth rate reflects the conservative end of publicly reported manufacturing expansion estimates; the Abundance Scenario applies a higher rate within this range
  - U.S. and PJM shares of future turbine deliveries held at 2025 levels, reflecting continued concentration of data center-driven demand.
  - Projected global turbine deliveries allocated to PJM based on geographic distribution of orders reflecting DC demand growth.

Interpretation

- Represents a physical upper bound on gas capacity that could be delivered to PJM.
  - Reflects turbine manufacturing capability, supply chain availability, and expected turbine allocation patterns.
  - ***Values do not represent a forecast of expected gas additions.***

Sources: Utility Dive GE Article; GE Vernova; S&P Global Article

Storage build limits (2026-2047) are assumed to grow and then plateau

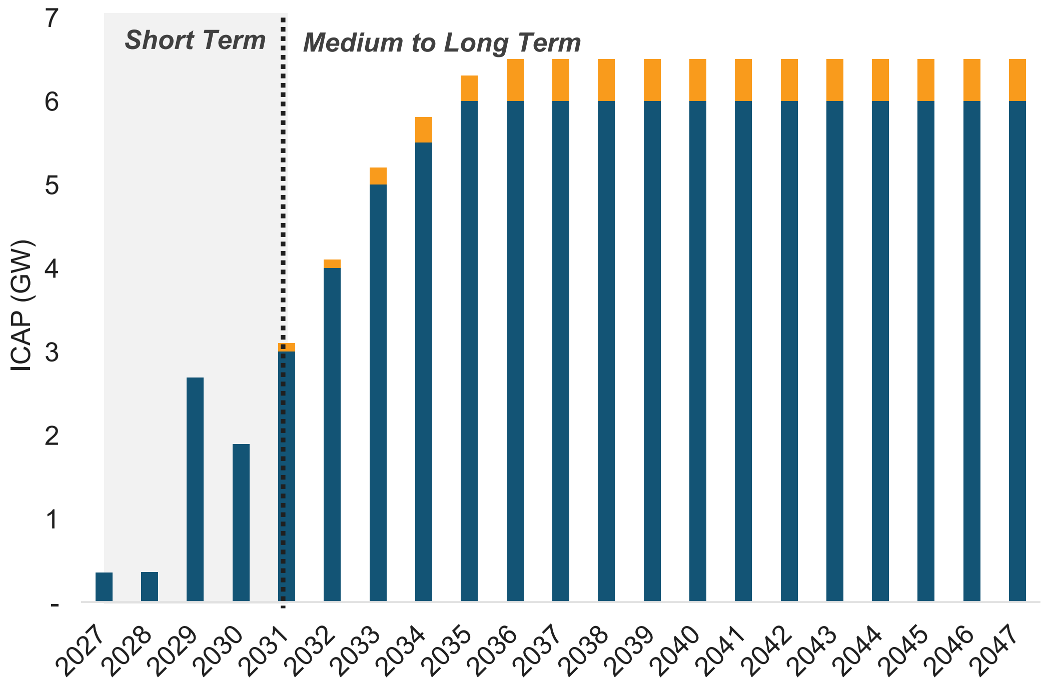

Methodology

Lithium Ion

- Near-term storage limits based on the active storage development pipeline and projects advancing toward commercial operation.
- Mid- to long-term additions ramp up to ~6 GW/year and remain flat, reflecting a realistic buildout pace, state policy goals and efficiency gains.

Other Technologies

- Storage limits are benchmarked to CAISO's technology adoption curve, which reflects a more mature market
- Non-lithium-ion storage is assumed to maintain the same share of total BESS additions observed in CAISO, applied with a five-year lag to reflect PJM's slower market development.
- The PJM lithium-ion buildout serves as the baseline forecast to determine this constraint.

Interpretation

- Represents a physical upper bound on battery storage capacity that could be delivered to PJM.
  - Reflects battery cell and transformer manufacturing capacity, supply chain availability, and expected allocation of storage to PJM.
  - ***Values do not represent a forecast of expected BESS additions.***

**Sources: EIA; US DOE; CPUC AB1373; CAISO 20 Year Transmission Outlook; Virginia H.B. 895*

Nuclear additions are driven by announcements and policy-drivers

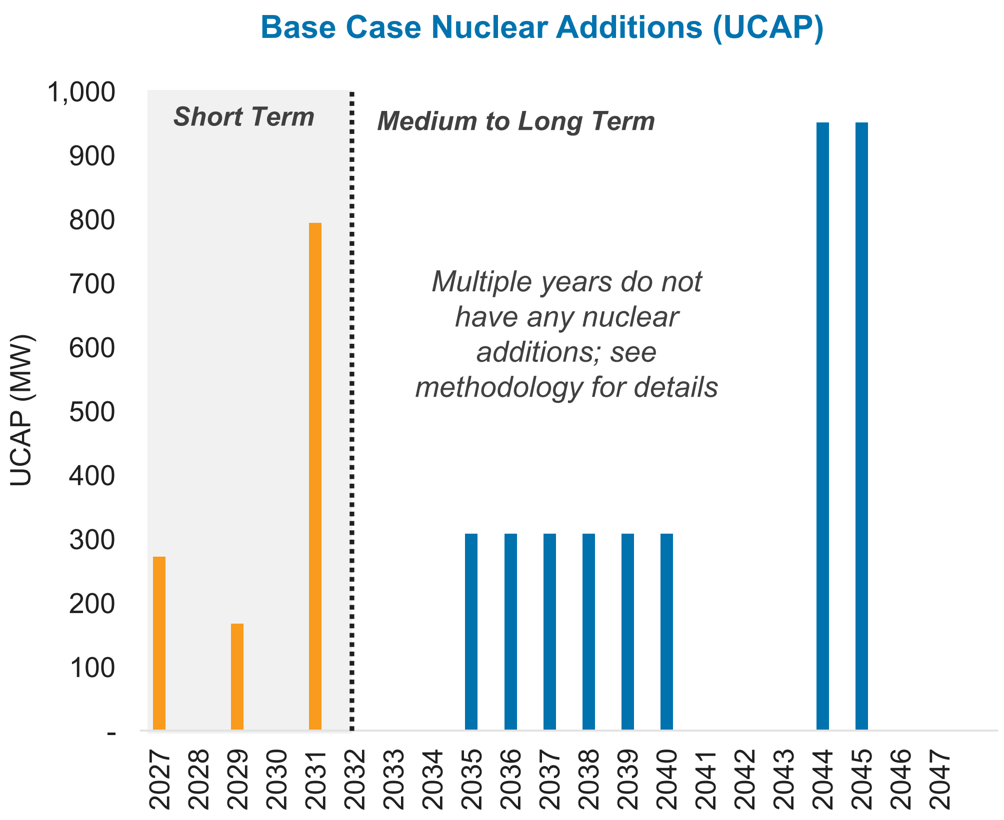

Sources	Additional Information
CRA Analysis	• PJM RRI and TC2 Queue • Dominion Energy 2025 IRP Update • Constellation FERC waiver filing (Mar 2026)

CRA analyzed announced and in-queue nuclear builds across a range of public sources for the near term, and drew long-term additions from utility IRPs and the LTCE economic buildout.

- Short term (2027-2030): nuclear additions are limited to uprates already in the PJM interconnection queue and RRI, plus one announced restart
  - Crane Clean Energy Center (835 MW restart) is the single largest near-term addition; the 2031 in-service date reflects the waiver submitted to FERC in March 2026
- Long term (2031-2047): nuclear additions do not come solely from the LTCE
  - Dominion: nuclear is included as an IRP addition - as a vertically integrated utility, Dominion can build and rate-base new nuclear, entering in the late 2030s
  - ComEd: new nuclear enters in the mid-2040s - CEJA-driven gas retirements leave the zone short of both capacity and energy after 2040, and the Aurora economic buildout selects nuclear to fill the gap

Near-Term Nuclear Uprates and Additions (in Queue)

Project	Zone	Type	MW Capacity	Online
Crane Clean Energy Center	PPL	Restart	835	2031
LaSalle	ComEd	Uprate	174	2027-30
Byron	ComEd	Uprate	43	2027-30
Braidwood	ComEd	Uprate	58	2027-30
Salem	PSEG	Uprate	228	2027-30

Existing supply in PJM

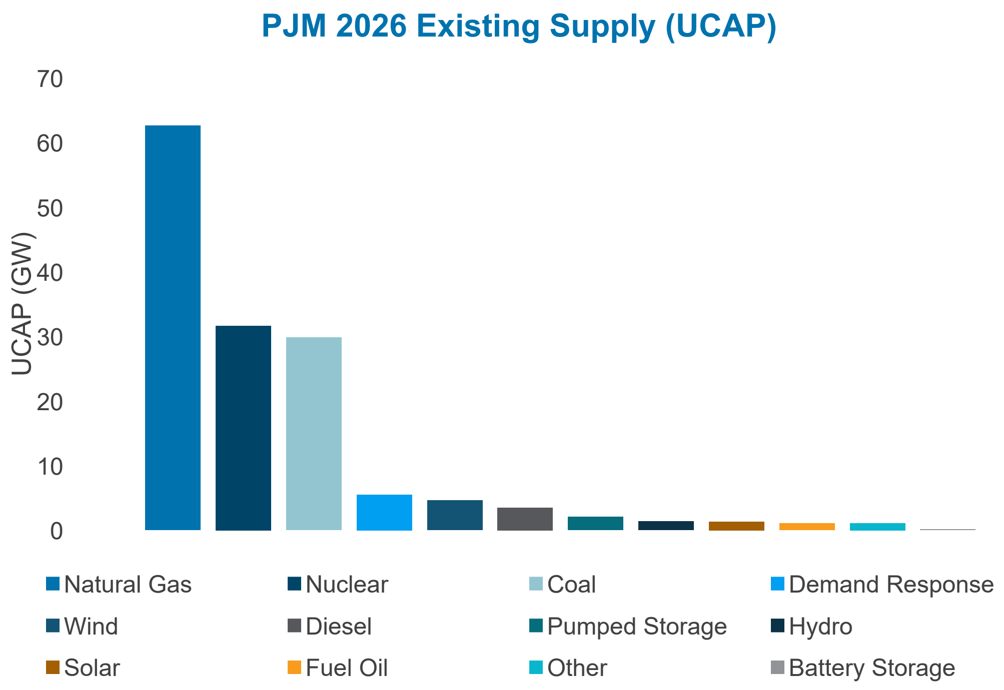

Sources	Additional Information
Energy Velocity	Used to find all tech types, except for gas, by zone (Compilation of EIA 860, Hitachi/ABB Research, NERC ES&D research)
PJM 2026 IMM Report	Source for gas ICAP by zone
ELCC Class Ratings	For 2025/2026 Base Residual Auction

Approach

- Total set of existing non-gas resources sourced from Energy Velocity (EV) by PJM zone
- ICAP of dispatchable generation (high ELCC) cross-referenced with PJM 2026 IMM Report
- ICAP converted to UCAP using 2026 ELCC class ratings

Tech type	ELCC Value (2026)	Notes
Gas	70%	Weighted average of CT and CC
Coal	83%	-
Nuclear	95%	-
BESS	52%	4-hour storage used

Existing fleet held at 2026 accreditation values as the baseline before annual ELCC evolution is applied

Interconnection queue in PJM

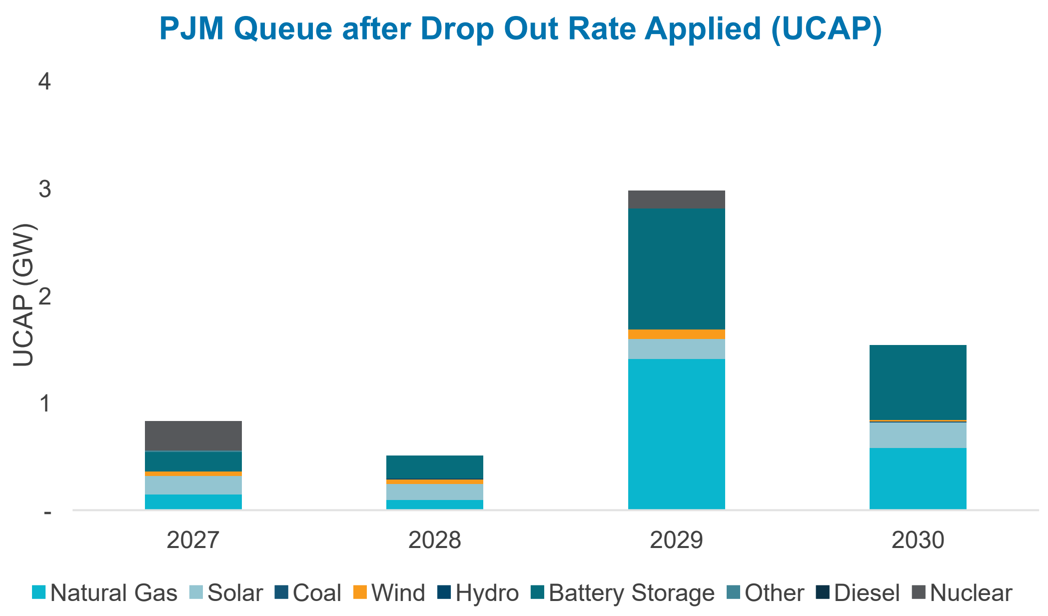

Sources	Additional Information
PJM Interconnection Queue	Legacy queue which includes older entrants. TC1 and TC2 resources excluded to avoid double-counting
TC1	Transition Cycle 1 projects, from PJM website
TC2	Transition Cycle 2 projects, from PJM website
IRPs	Gas projects included in IRPs automatically included in buildout. IRPs used include Dominion and APS.

Queue drives near-term supply; LTCE drives long-term supply. The queue provides an independent supply signal for 2027-2030.

Approach

- Assumed near-term additions (2027-2030) taken from PJM interconnection queue to avoid circularity between load forecast and supply forecast*
- Legacy queue, TC1, and TC2 combined; TC1/TC2 excluded from legacy count to avoid double-counting
- Gas additions limited to named projects from integrated utility resource plans (IRPs) (e.g., Dominion, APS), under construction, or an uprate from PJM's Reliability Resource Initiative (RRI)

Assumptions

- Only resources currently active in the queue included; "partially in-service" excluded
- TC2 projects without a COD distributed using TC1's tech-specific COD distribution
- Multi-fuel projects split across listed technologies
- Assumed future drop-out rates are applied by technology and are specific to each queue cohort (Legacy, TC1, TC2), reflecting the different attrition observed in each

Technology	Legacy queue [1]	TC1	TC2
Solar	69%	10%	55%
Wind	40%	10%	52%
Battery Storage [2]	69%	10%	55%
Source	*Historic queue data*	*BRA-Anchored*	*TC1 drop-outs*

** CRA recognizes some projects withdrew from the Transition Cycles and reapplied to Cycle 1. These projects are captured as generic units in the LTCE buildout, unless specifically named in an IRP.*

[1] Average from 2027 – 2030.

[2] Battery Storage drop-out rate is set equal to that of solar due to limited historic data.

CRA Charles River Associates

How CRA handles Cycle 1

PJM Cycle 1 Total Submissions vs. Current Queue (Nameplate Capacity)

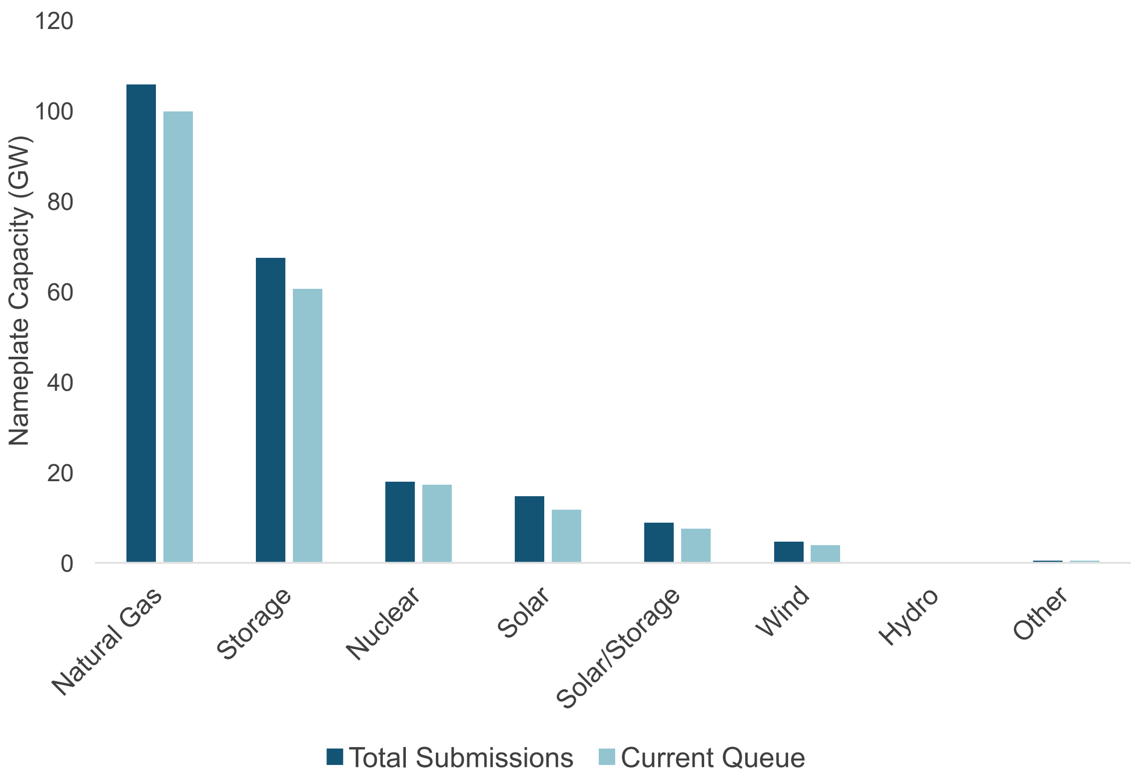

CRA incorporates Cycle 1 into the LTCE as generic additions, rather than as dated queue entries, for the base case

- Additional Cycle 1 projects are captured as generic units, realized post-2030, under the Abundance Scenario to acknowledge interconnection queue improvements

Why?

- **Scale:** 711 projects, representing 201 GW of nameplate capacity, are currently active in Cycle 1
- **No reliable timing or attrition signal:** Cycle 1 is the first cycle under the reformed process, so its new structure and shifting dynamics make it difficult to reliably extrapolate CODs or attrition rates from the Transition Cycles
- **Timing points to the LTCE window:** New queue processing takes ~2 years, combined with multi-year build times for high-ELCC resources (gas, nuclear), the earliest realistic online date falls within the LTCE period (2031+)

LTCE results show a steady build of resources over time across all technology types

PJM Supply Mix (UCAP)

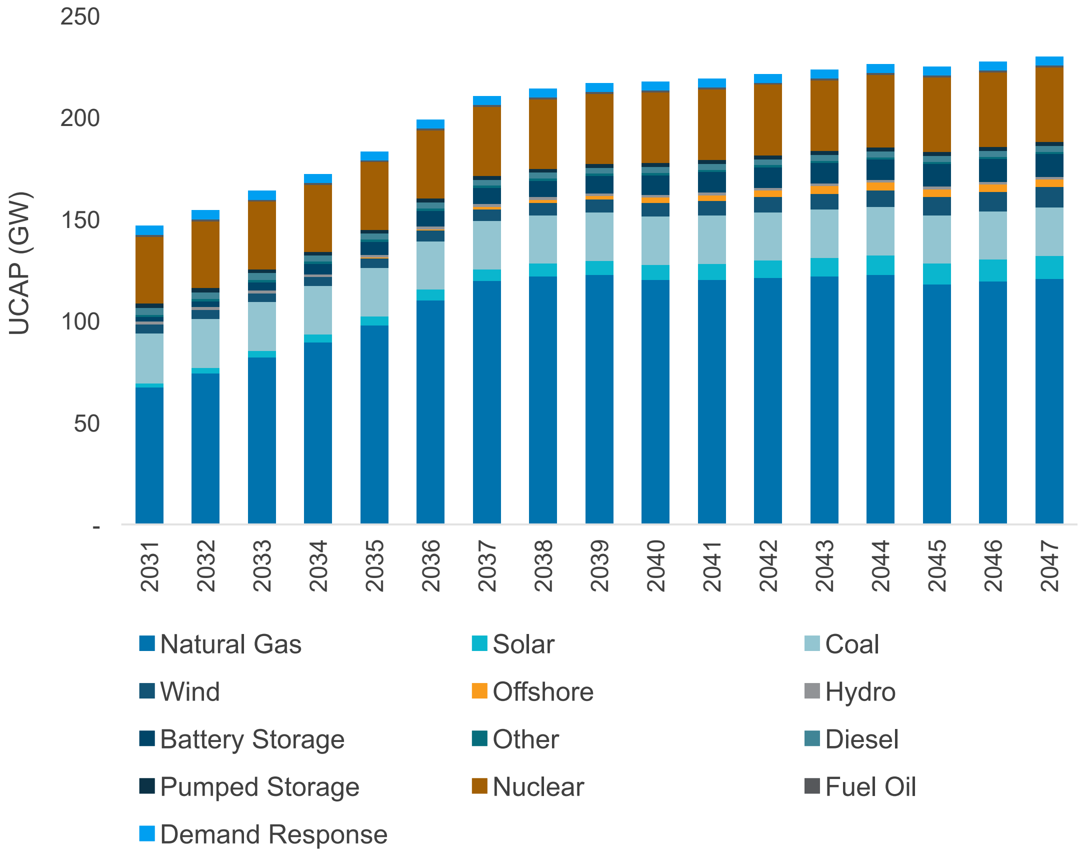

Approach

- Long-term supply (2031-2047) from CRA's Aurora LTCE model, with endogenous build and retirement decisions
- LTCE picks up where the near-term queue leaves off, subject to build limits and transmission limits
- Cycle 1 projects represented through generic LTCE additions rather than dated queue entries, given queue processing time *(see Cycle 1 slide)*

Assumptions

LTCE results reflect:

- **Fuel prices** transition from near-term market forwards to CRA long-term fundamental forecasts.
- **Policy reflected:** OBBBA, CEJA, state RPS, and Clean Grid Act; new-build costs from NREL/EIA; ITC/PTC incorporated.
- **Environmental compliance** costs applied (CO2 via RGGI and applicable programs, SO2, NOx).
- **PJM-published ELCC accreditation** applied to all resource types
- **Gas and storage builds** constrained to annual build limits; any capacity above the limit is pushed to later years until it fits.
- **Post-2038, no power or transmission constraint is applied; gas and storage limited only by the build limit.**

Sources	Additional Information
CRA Analysis	Long-term supply forecast based on LTCE model results, including endogenous build and retirement decisions.

CRA Charles River Associates

Modest retirement of existing plants

PJM Announced & Economic Retirements (UCAP)

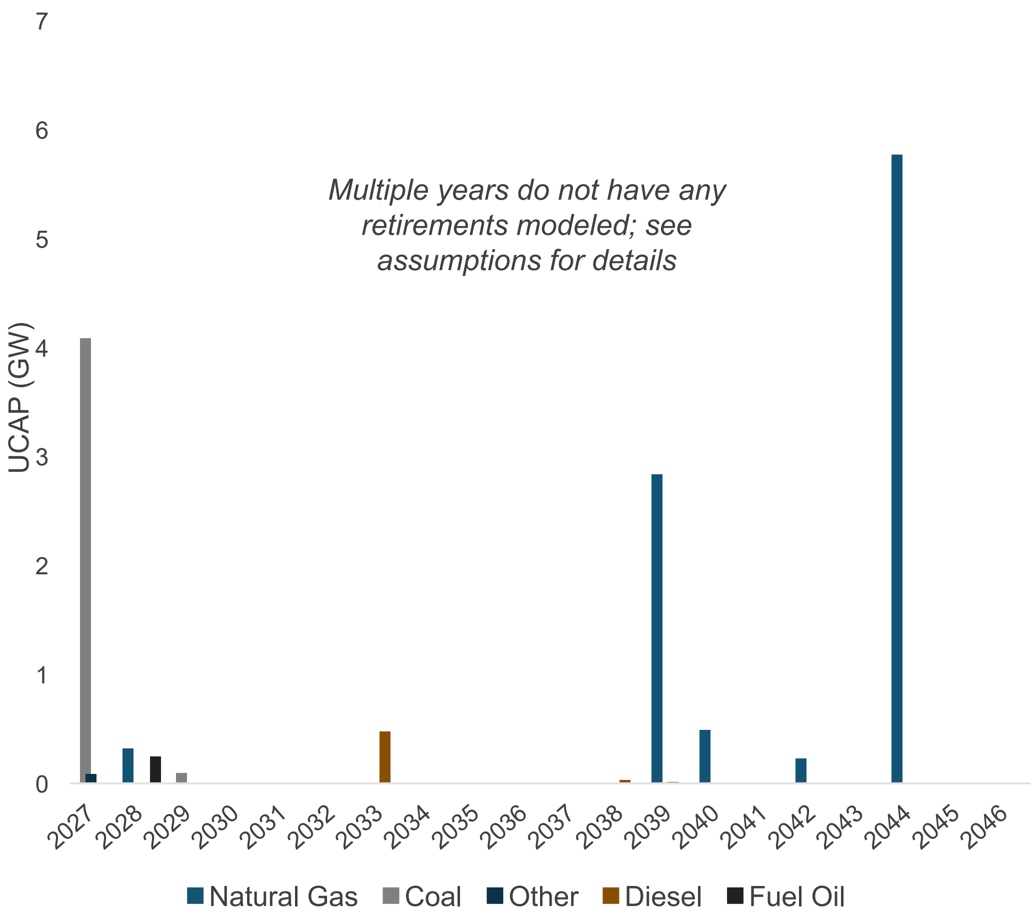

Approach
Two retirement channels:
- **Short term:** Announced retirements as fixed inputs
- **Long term:** Economic retirements determined endogenously by the LTCE

Assumptions

Short-term retirements
- Based entirely on public deactivation requests sourced from Energy Velocity
  - Coal retirements concentrated in 2027-2028 (~6.1 GW)

Long-term retirements
- Post-2035, retirements are driven by state/federal policy (e.g., CEJA) or revenue-adequacy challenges
- Economic retirements are a model output, not an input, and respond to prices, fuel costs, and accreditation

Sources	Additional Information
Energy Velocity	Used to find all tech types, by zone (Compilation of EIA 860, Hitachi/ABB Research, NERC ES&D research)
Public Policy	Legislated retirement schedules and emissions policy (e.g., Illinois CEJA, RGGI, EPA rules) informing long-term deactivations

Capacity accreditation is taken from PJM-published values

ELCC Values from PJM BRA and PJM Preliminary Class Ratings

Year	Gas	Solar	Coal	Wind	OSW	Hydro	Battery Storage	Diesel	Pumped Hydro	Nuclear	Fuel Oil	DR
2026	0.70	0.09	0.83	0.41	0.68	0.40	0.52	0.91	0.4	0.95	0.61	0.85
2027	0.70	0.09	0.83	0.41	0.68	0.40	0.52	0.91	0.4	0.95	0.61	0.85
2028	0.75	0.07	0.85	0.34	0.6	0.35	0.59	0.93	0.37	0.96	0.61	0.91
2029	0.71	0.07	0.83	0.35	0.57	0.37	0.42	0.91	0.37	0.95	0.62	0.77
2030	0.71	0.07	0.83	0.34	0.53	0.38	0.37	0.91	0.38	0.95	0.63	0.74
2031	0.72	0.07	0.82	0.31	0.47	0.39	0.30	0.91	0.39	0.95	0.64	0.7
2032	0.74	0.07	0.81	0.26	0.38	0.39	0.25	0.91	0.39	0.95	0.66	0.68
2033	0.75	0.07	0.81	0.22	0.32	0.38	0.25	0.91	0.38	0.96	0.68	0.68
2034	0.75	0.07	0.80	0.20	0.28	0.38	0.24	0.91	0.38	0.95	0.69	0.67
2035*	0.76	0.07	0.80	0.19	0.26	0.38	0.23	0.91	0.38	0.95	0.70	0.66

Sources	Additional Information
BRA ELCC Class Rating	Final class ratings from published Base Residual Auction planning parameters
PJM Preliminary ELCC Class Ratings (2027/28 – 2035/36)	PJM-published preliminary class ratings for delivery years without final BRA results

** ELCCs held constant post 2035/36 DY*

Assumptions

- All ELCC values are taken directly from PJM published class ratings; CRA did not independently derive or adjust any accreditation values
  - Published BRA class ratings used for DYs with auction results
  - PJM Preliminary Class Ratings (2027/28-2035/36) used for DYs that do not yet have auction results or published planning parameters
- Gas ELCC applied as a capacity-weighted average of PJM's published CC and CT class ratings, reflecting PJM's fleet composition
- 4-hour duration used as the storage class rating across the forecast
- **ELCC values held constant post 2035/36 DY**

Charles River Associates

Existing supply in other regions

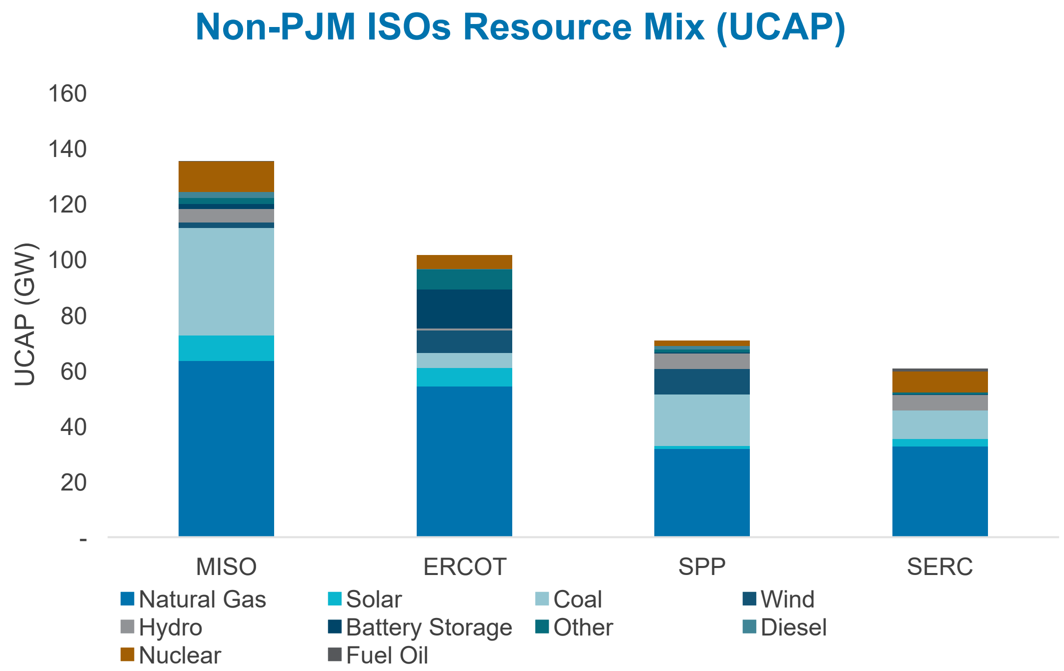

Sources	Additional Information
Energy Velocity	Used to identify all resources, by ISO (Compilation of EIA 860, Hitachi/ABB Research, NERC ES&D research)
ISO-Specific Accreditation Values	Most recent accreditation values used to convert ICAP to UCAP

Approach

- Same method as PJM: 2026 existing fleet in ICAP by region (MISO, SPP, ERCOT, SERC) from Energy Velocity
- ICAP converted to UCAP using region-specific accreditation values

Assumptions

- Regions selected as the primary candidates for data center migration out of PJM
- Only resources listed as in operation are included
- If accreditation values are not available for a resource in a specific ISO, the value is based on a neighboring ISO's

Variable	Accreditation Value
Gas	89 – 92%
BESS	61 – 100%
Coal	72 – 90%
Nuclear	95 – 98%

Interconnection queue in other regions

Non-PJM Queue Additions (UCAP, GW)

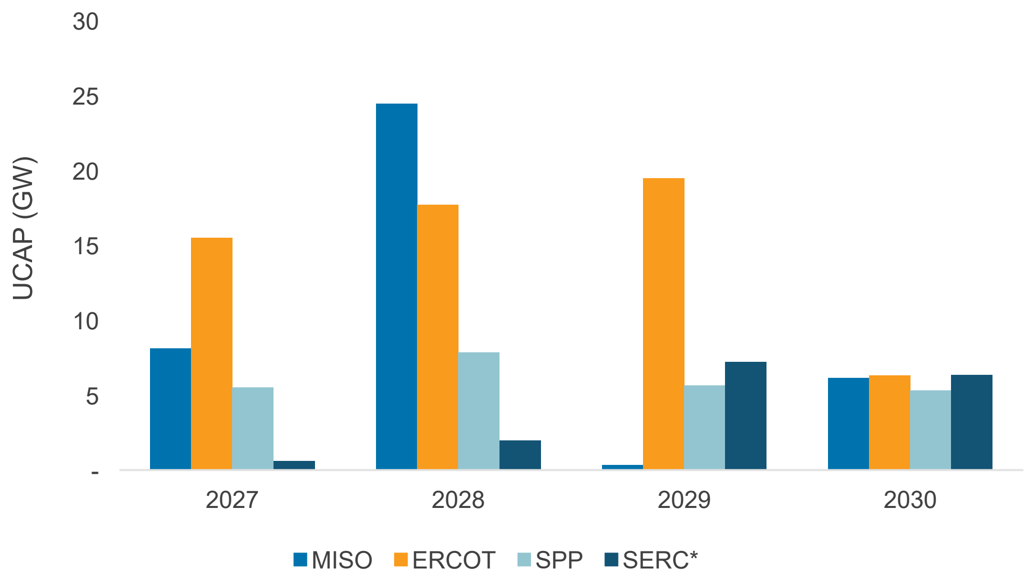

Sources	Additional Information
MISO	MISO generation queue, MISO ERAs Queue
SPP	SPP generation queue
ERCOT	ERCOT generation queue
SERC*	Utility IRPs and CRA LTCE Analysis

Approach

- Near-term additions (2027-2030) driven by the interconnection queue rather than LTCE, when possible, to avoid circularity
- ICAP converted to UCAP using region-specific accreditation values

Assumptions

- Drop-out rates are tech-specific and ISO-specific, based on historic averages. Nuances include:
  - ERCOT: historic withdrawal rates; changing ISO environment results in required Interconnection Agreement to be counted additions in 2027-28 (this constraint loosens in '29-30)
  - MISO: a 10% future drop-out rate is applied to Expedited Resource Addition Study (ERAS) gas resources in lieu of a historically-informed drop-out rate

**Note – From CRA LTCE Analysis*

CRA projects modest additional headroom in non-PJM regions

Non-PJM Served DC Load and Available Headroom

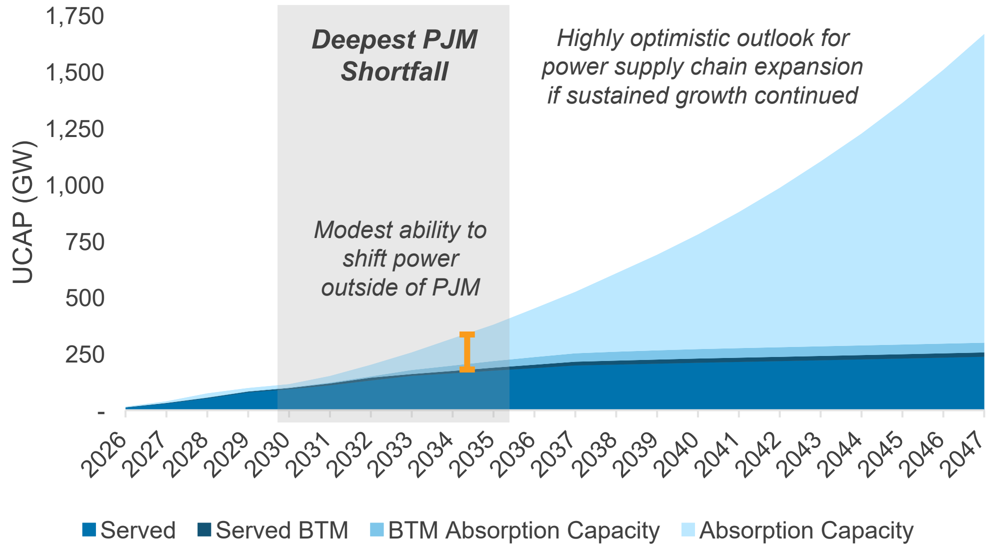

Sources	Additional Information
MISO	MISO queue, CRA LTCE Analysis, MISO DC Forecast
SPP	SPP queue, CRA LTCE Analysis, SPP DC Forecast
ERCOT	ERCOT queue, CRA LTCE Analysis, ERCOT DC Forecast
SERC*	Utility IRPs and CRA LTCE Analysis, CRA DC Forecast

**Note – No official DC forecast available; CRA compiled internal forecast*

Approach

- Near-term additions (2027-2030) follow queue methodology
- Long-term LTCE adjusted to meet build limit constraints for natural gas and storage to account for all potential generation resources
- Data center load forecasts held at ISO's internal view*
- Build limits and data center forecasts implicitly simulate the attractiveness to data centers in these regions and the impact of policy on the ability to develop new generation.

Assumptions

- If regional forecast ends prior to 2047, DC demand is extrapolated from current data and is assumed to track growth trends of ISOs with available data

Absorption capacity = max (0, UCAP – forecasted peak load*(1+RM)

- ❖ Measured as the ability of an ISO or region to serve additional DC load than what is in their forecast; **a positive value indicates room to host relocated PJM load**

BTM Absorption capacity = max (0, BTM capacity - unserved DC load*(1+RM)

- ❖ Measured as the ability of an ISO or region to host additional DC load than what is in their forecast BTM; **a positive value indicates room to host relocated PJM load**

- Output feeds the "where does displaced PJM load go" question

CRA Charles River Associates

Meeting agenda

Topic	Timing
Introductions and July 13 LAS Meeting Recap	9:10 – 9:20 AM
Top-Down: Unconstrained View	9:20 – 9:35 AM
Top-Down: Constraints	9:35 – 10:00 AM
Bottom-Up Forecast	**10:00 – 10:15 AM**
Load Profile Assumptions	10:15 – 10:25 AM
Forecast Reconciliation and BTM	10:25 – 10:45 AM
Scenarios	10:45 – 11:00 AM
Break	*11:00 – 11:10 AM*
Q&A	11:10 – 11:55 AM
Next Steps & Close	11:55 – 12:00 PM

To determine where projects are most likely to be sited, each PJM zone was assessed against key DC developer siting criteria

Zonal Siting Criteria
High Priority
These factors are high priorities as developers seek to obtain the primary objective of "speed to market."

Zonal Siting Criteria
Nice to Have
These factors are mostly seen as pluses rather than requirements.

- Power Availability ← *Covered elsewhere*
- State / Local Regulatory Support (headwind, neutral, tailwind)
- Legacy Market Presence
- Site Development Cost & Feasibility
- Resilience / Physical Risk Profile
- Fiber Connectivity
- Water
- Low Carbon Power Access ↑ *Uniform across PJM*

Evaluation Criteria

For each PJM Zone rank Relative Favorability
1 = Low
2 = Average/Medium
3 = High

X

Assign Category Importance Multiplier
Higher weighting for "High Priority" vs. "Nice to Have"; Scale of 1-5

= Zonal Evaluation Criteria

Targeted Outcome:
Assessment of the relative attractiveness of PJM zones to weight pipeline projects and allocate the top-down forecast, with growth concentrated in higher-scoring zones and less weight assigned to projects in lower-scoring zones

CRA's model treats market formation as a two-part process: first, whether a dormant zone becomes a durable market; second, how annual demand is allocated once a zone is live

What Gets Built In 'Live' Markets vs. What Takes Off

Core finding: CRA observed multimodal behavior in the historical data: existing markets build proportionally to existing market share but are capped based on physical limits. New markets take off proportional to national growth. Attractive markets are more likely to take off but are hard to predict beforehand.

'Live' Markets

How much does each 'live' market build?

Growth follows installed scale + discounted queue, capped by available power. CRA includes the expected realized pipeline in existing share.

New/ Emerging Markets

Which dormant market launches next?
Informed by zonal attractiveness score (next slide)

How large is the launch?
Informed by national demand growth

How fast can each market grow?
Informed by historical growth tiers; caps proxy permitting, labor, and other local constraints

CRA conducted detailed research to assign each criterion a favorability score, then applied an importance multiplier to reflect its relative influence on DC siting decisions

Category	Research Conducted	1 = Lower Favorability	2 = Medium Favorability	3 = High Favorability	X	Category Importance Multiplier 1 = Lowest, 5 = Highest
State / Local Regulatory Support	Existing/proposed State and local bills that would hinder DC development mapped to PJM zones	Restrictive or uncertain approval environment.	Permitting path exists but with some risk.	Supportive policy and clear approval path.		4
Legacy Market Presence	Sum of commissioned, available, under construction and planned MW of power as per merged datacenterHawk + Cleanview data set	Limited existing presence or customer activity.	Emerging presence, but pipeline depth or customer momentum remains uncertain.	Established market with proven delivery track record and active customer demand.		3
Site Development Cost & Feasibility	• **Energy costs:** PJM historical zonal hourly LMP (Energy Velocity) • **Land costs:** Rural farm real estate $/acre by State (USDA 2025) • **Tax incentives:** State-level DC tax incentives research	Limited or uncertain incentives, excess vacancies in colocation availability.	Some incentives available, but not clearly differentiating, average vacancies.	Strong incentives that improve site competitiveness, low vacancies.		2
Resilience / Physical Risk Profile	FEMA National Risk Index v1.20 (December 2025) by county mapped to state and zone; includes Inland Flooding, Severe Storms, Wildfire, Winter Storms, Extreme Heat risks	High physical risk exposure.	Manageable risk with mitigation.	Low risk and strong resilience profile.		2

CRA Independent Assessment of Project Siting Realization Probabilities

CRA conducted detailed research to assign each criterion a favorability score, then applied an importance multiplier to reflect its relative influence on DC siting decisions *(Continued)*

Category	Research Conducted	1 = Lower Favorability	2 = Medium Favorability	3 = High Favorability X	Category Importance Multiplier 1 = Lowest, 5 = Highest
Fiber Connectivity	ITU Fiber Optic Maps (high-level and relative view evaluated between PJM zone)	Limited connectivity or redundancy.	Connectivity available but may require buildout.	Strong fiber access, redundancy, and latency profile.	2
Water	World Resources Institute's Aqueduct 4.0 Water Risk Atlas database	Material water or labor constraint.	Available but with some cost or planning risk.	Adequate water and labor to support growth.	1
Low Carbon Power	Treated as uniform across PJM	N/A	N/A	N/A	1

CRA Independent Assessment of Project Siting Realization Probabilities

CRA Charles River Associates

New projects are ‘selected’ to achieve grid connection in proportion with existing market share

Statistical analysis showed that year-over-year capacity additions are highly correlated with existing market share, but growth gradually spreads into new markets over time. This reflects data center developers' preference for expanding in established markets while simultaneously pursuing development opportunities in smaller and emerging markets.

Effective installed base

Each zone's weight starts from operating grid capacity plus behind-the-meter capacity plus the projects it is expected to complete from the queue.

Risk-weighted queue

Queued projects count only in proportion to how likely they are to be built. The older a project is, the more likely it is to drop out.

Share raised to a power (<1)

Effective shares are scaled by an exponent that sharpens the advantage of established zones, and that value acts as a floor no zone falls below.

Dormant zones stay in play

Zones with no capacity yet receive a small collective allocation, added on top so leading zones never lose ground. If existing active zones can’t keep up, new zones may ‘ignite’.

Net Effect: **Growth concentrates where capacity and credible pipelines already exist, while emerging zones keep a realistic chance of landing the next project and capacity gradually spreads across the market.**

Based on CRA's assessment, emerging zones are classified by their relative potential to ignite and sustain data-center growth

Live Markets (>300 MW)	High-Potential Emerging	Emerging	Measured growth
Established data-center markets where future growth depends primarily on existing scale, pipeline, and available power—not overall attractiveness rankings.	*Dormant or limited-scale markets with the strongest overall conditions to become the next durable growth markets.*	*Markets with credible growth potential but one or more constraints that may limit the timing or scale of development.*	*Markets facing multiple or material disadvantages that reduce the likelihood of significant near-term growth, relative to other PJM regions.*
• Dominion • PSEG • ComEd • AEP • APS • PPL	• ATSI • DUQ • DAYTON (strong pipeline)	• PECO • DEOK • EKPC • METED • DPL • BGE	• PENELEC • AECO • PEPCO • JCPL • RECO
Feedback welcome! Results are subject to change on receipt of verifiable pipeline and policy information.	Assign the highest ignition probabilities and realization weightings among non-live markets ☆ ***Markets Most Likely to Ignite***	Apply zone-specific ignition probabilities and moderate realization weightings.	Assign lower ignition probabilities and realization weightings unless project-specific evidence supports an adjustment.

Readiness tiers are used to evaluate the maturity of the pipeline

Tier 1: Early / Speculative
Limited evidence of near-term execution, including speculative requests, uncertain financing, smaller developers, or less durable use cases.

→ Attrition

Tier 2: Advancing / Credible
Clearer development momentum, including construction activity, larger load requests, known developers, colocators, or stronger financial indicators.

→ Attrition

Tier 3: Materialized Load
Strongest evidence of materialization, including commissioned power, leased space, nearby generation, established developers, or public financial backing.

→ Attrition

How the queue is modeled

- Projects are assigned to readiness tiers based on observable development indicators.
- A **hazard function** is applied to estimate the probability that projects advance, remain in queue, or fall out over time.
- These probabilities are learned from observed project progression data rather than assumed uniformly across the queue.
- “Generic” projects are added separately to preserve expected future demand in areas where known queues may not fully capture market activity.
- If a zone shows signs of acceleration, generic demand allows the forecast to reflect emerging opportunity before named projects appear in the queue.

CRA groups nearby facilities from the same developer into campuses to jointly simulate buildings that operate as a shared investment decision

1. **PJM Scope** → The analysis is restricted to PJM facilities, with hyperscale and colocation projects evaluated separately

2. **Provider Segmentation** → Facilities are evaluated only against other sites from the same company/provider, preventing projects from different developers from being combined into a single campus

3. **Geographic Clustering** → Facilities within the same company/provider are grouped based on latitude and longitude using a ***2-mile geographic threshold***, with isolated facilities retained as single-building campuses

4. **Cluster Validation** → The full geographic span of each proposed campus is validated, and overextended chained clusters are split so that no final campus exceeds the 2-mile threshold

5. **Campus Assignment** → Each facility is assigned a final Campus ID, creating an External ID-to-Campus ID crosswalk that supports campus-level aggregation, sequencing, and ramping in the bottom-up forecast

Source: DatacenterHawk; CRA Analysis

PJM is altering its approach to annual forecasting by introducing CRA's independent forecast to improve the evaluation of large loads submitted by LSEs/EDCs for the 2027 forecast

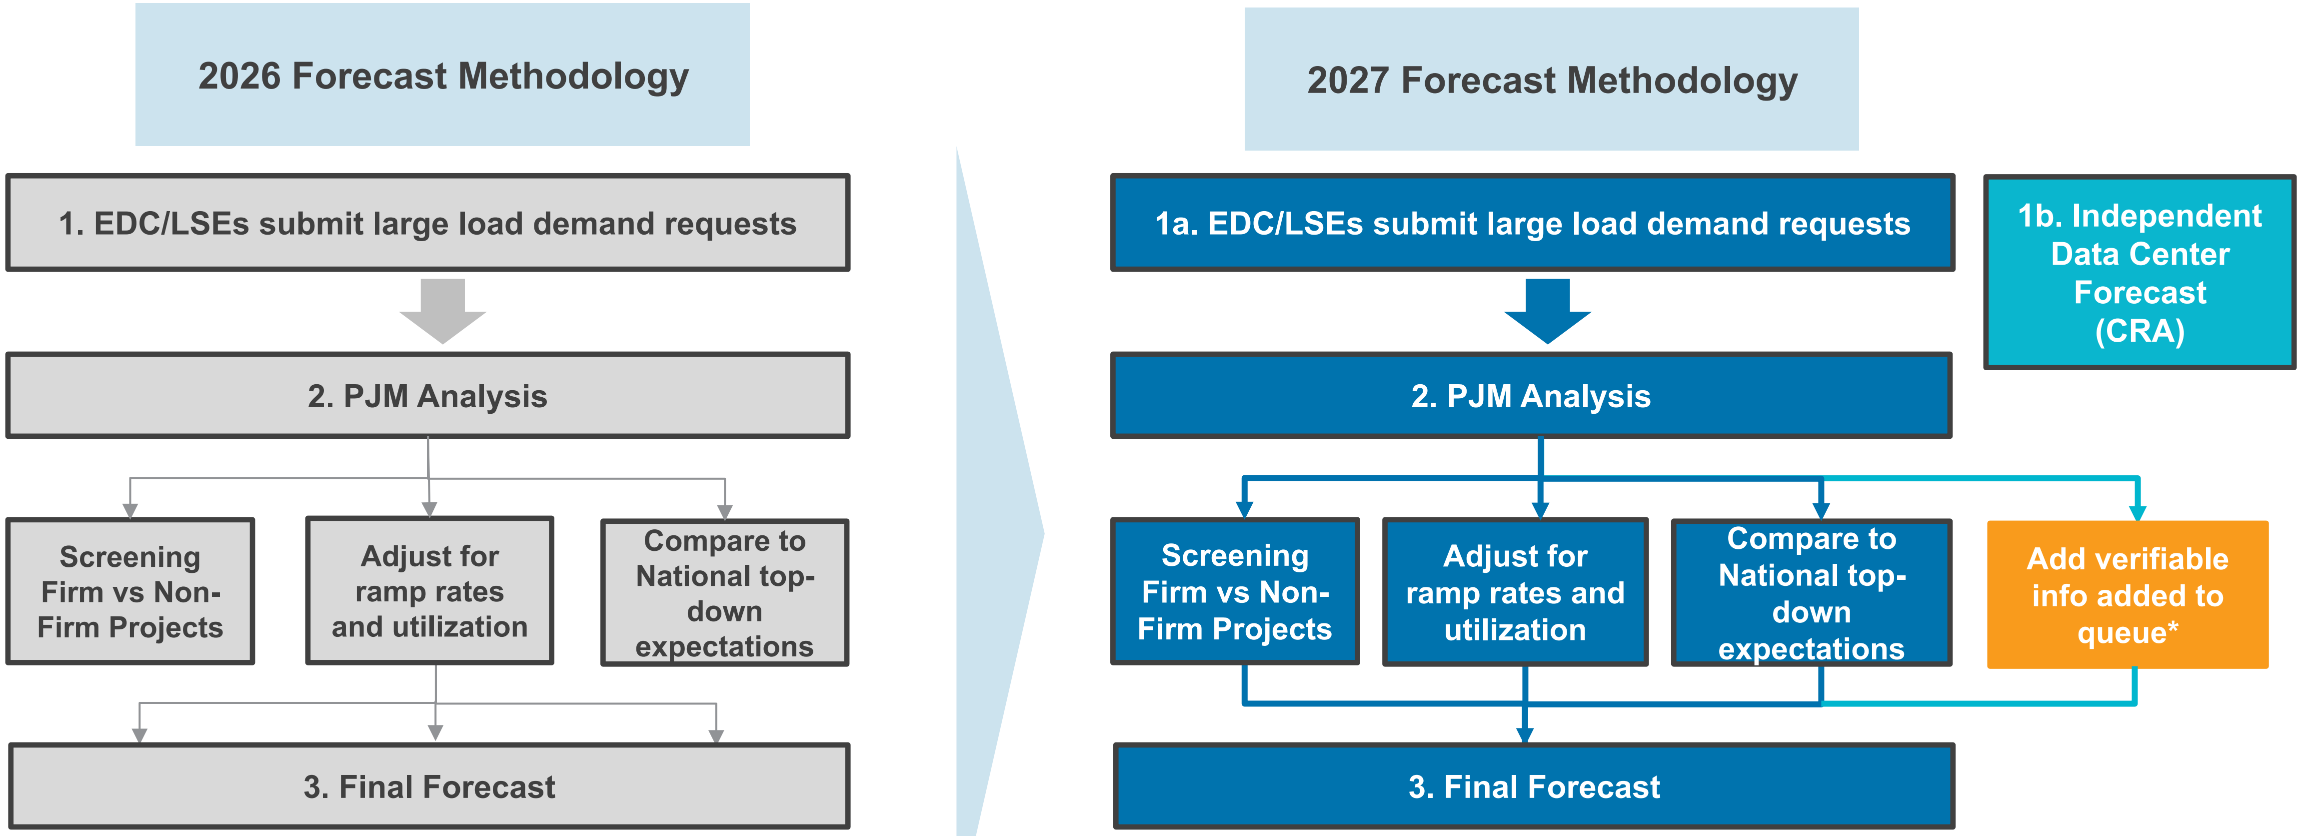

**CRA will incorporate project-specific information or adjust applicable build limits where stakeholders provide sufficient, verifiable evidence to support the change.*

Meeting agenda

Topic	Timing
Introductions and July 13 LAS Meeting Recap	9:10 – 9:20 AM
Top-Down: Unconstrained View	9:20 – 9:35 AM
Top-Down: Constraints	9:35 – 10:00 AM
Bottom-Up Forecast	10:00 – 10:15 AM
Load Profile Assumptions	**10:15 – 10:25 AM**
Forecast Reconciliation and BTM	10:25 – 10:45 AM
Scenarios	10:45 – 11:00 AM
Break	*11:00– 11:10 AM*
Q&A	11:10 – 11:55 AM
Next Steps & Close	11:55 – 12:00 PM

Characteristic hourly shapes

Public data on the operational load profiles and load factors of data centers is limited, and potential evolution of load profile and load factors is uncertain. LBNL's data center load shape generator will be used as a foundation to create the characteristic load shape.

Approach

- Characteristic hourly shapes will be anchored to LBNL's publicly available Data Center and Industrial Electrical Load Shape Generator
- Any changes based on expert conversations will be documented

Example of LBNL's Characteristic Hourly Shape for Business Diurnal Hourly Load

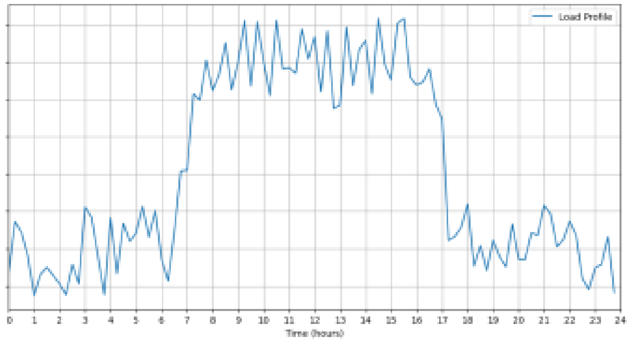

Example of LBNL's Characteristic Hourly Shape for Flat Hourly Load

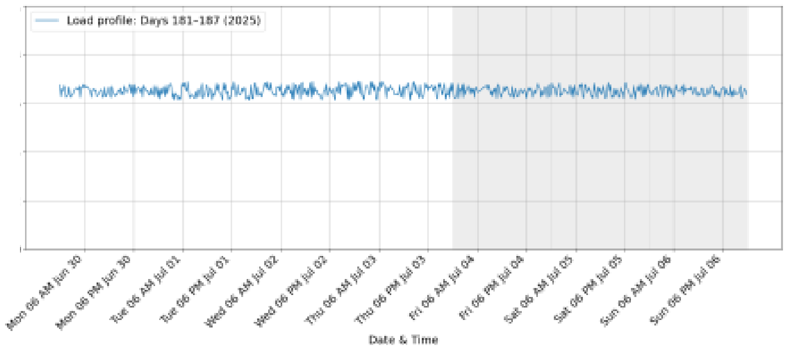

Identify Characteristic Shapes

To best capture 'typical' hourly behavior for different data center workloads (normalized to 0~1)

Incorporate Stochasticity

To best capture a range of potential events, rather than a single characteristic shape

Apply Forecasted Capacity

Scale normalized by expected data center growth by zone, distributed by estimated relative compute needs for

Source: LBNL Data Center and Industrial Electrical Load Shape Generator

Note: CRA acknowledges that the exact evolution of hourly profiles is uncertain and is rapidly evolving. LBNL provides a solid, publicly verifiable foundation to create characteristic hourly load. Any modifications to LBNL's outputs will be documented.

Ramp rates (absent constraints)

Larger data center campuses generally ramp with higher MWs, but are proportionally slower relative to smaller sites

Larger Data Campuses Ramp with Bigger MW Increments

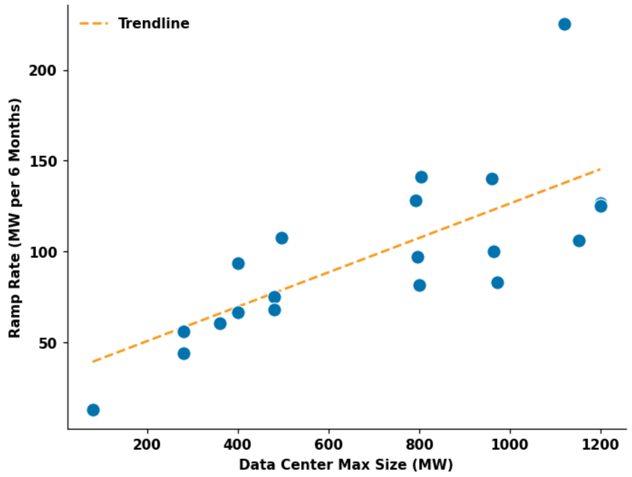

…but Slower as a Share of Their Own Maximum Capacity

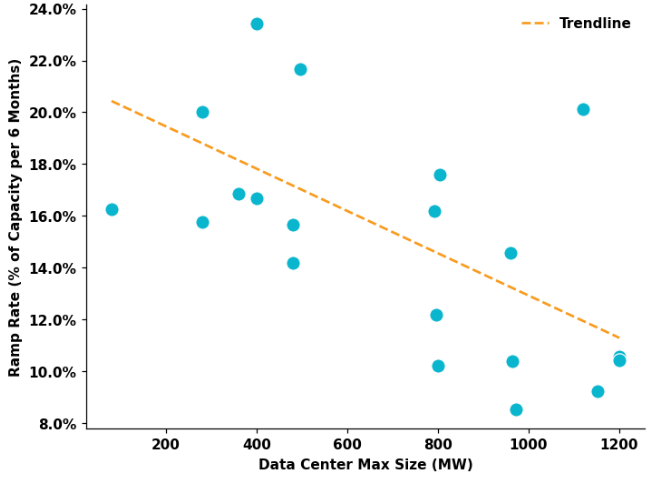

Note: calculated as the change in observed MW capacity between each month and the same data center's capacity 6 months earlier, averaged across the ramp period (which is between the first MWs coming online, and when the data center's MWs are equal to the maximum observed in the data)

Approach

- **Learn the pattern:** Fit historical ramp rates and timing as a function of data center size.
- **Preserve uncertainty:** Capture how different campuses deviate from the size-based trend.
- **Generate realistic ramps:** Combine the expected size-based curve with observed deviations from similar projects to create plausible MW ramp trajectories

Example of modeled data center ramps: larger data centers ramp slower than smaller data centers

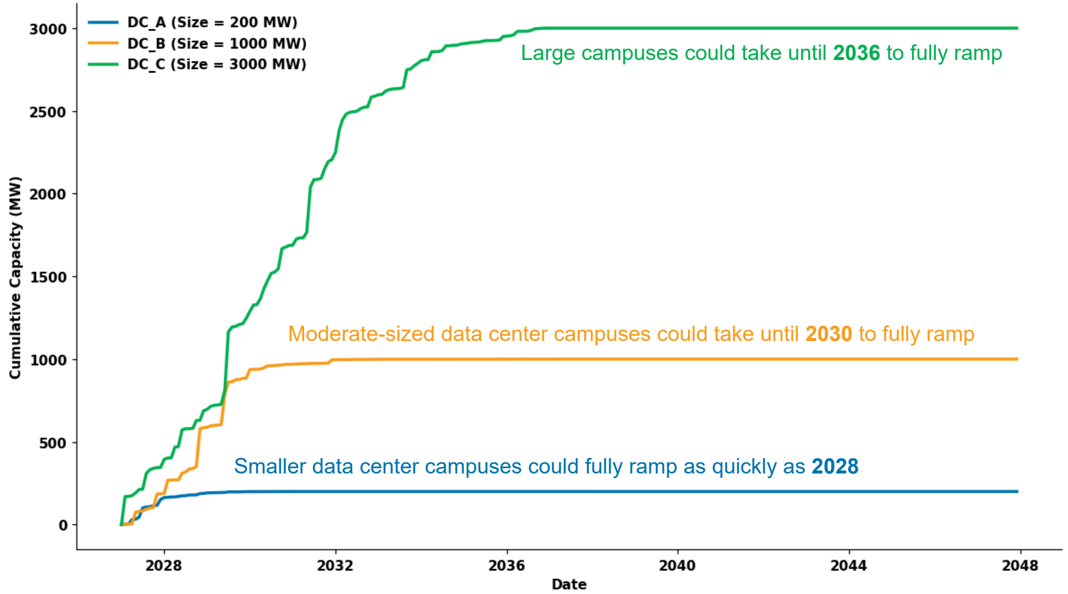

Example where all data center campuses start their ramp at the same time.

Assumptions

- Assumes size is the key dominant factor when considering data center ramp rates – the larger the load, the slower the ramp to full capacity
- This is to ensure that CRA is not selecting ramp rates based on other business-specific characteristics to preserve objectivity

Source: LSE Submissions

IT load is typically a fraction of total metered demand, as data centers typically require non-IT supporting equipment to operate effectively

“Total Metered Demand”

“Non-IT Load”

May include (but not limited to) the power needed by:

- Transfer switches
- Upstream transformers
- Uninterruptible Power Supply (UPS)
- Cooling towers
- Computer room AC/AH
- Fire-suppression systems
- Water-detection systems
- Building management systems

“IT Load”

May include (but not limited to) the power needed by:

- Servers
- Storage equipment
- Network equipment
- Telecommunications equipment
- Internal device power supplies
- Internal computing components
- Cooling devices that are integral to, and powered through the IT equipment

Note: Illustrative and informational purposes, not necessarily drawn to scale

Power Usage Effectiveness (PUE) is the ratio between Total Metered Demand and IT Load

$$PUE = \frac{Total\ Metered\ Demand}{IT\ Load}$$

Approach

- Apply PUE to IT load to estimate Total Metered Demand

$$\underbrace{Total\ Metered\ Demand}_{\text{CRA Estimates}} = \overbrace{PUE}^{\text{EPRI}} * \underbrace{IT\ Load}_{\text{datacenterHawk}}$$

Assumptions

- PUE estimates for data centers at 1.32 per EPRI’s capacity-weighted averages for facility-level data (link)

Sources: Recommendations for Measuring and Reporting Overall Data Center Efficiency, EPRI

Meeting agenda

Topic	Timing
Introductions and July 13 LAS Meeting Recap	9:10 – 9:20 AM
Top-Down: Unconstrained View	9:20 – 9:35 AM
Top-Down: Constraints	9:35 – 10:00 AM
Bottom-Up Forecast	10:00 – 10:15 AM
Load Profile Assumptions	10:15 – 10:25 AM
Forecast Reconciliation and BTM	**10:25 – 10:45 AM**
Scenarios	10:45 – 11:00 AM
Break	*11:00– 11:10 AM*
Q&A	11:10 – 11:55 AM
Next Steps & Close	11:55 – 12:00 PM

Methodology overview to produce 20-year hourly load forecast

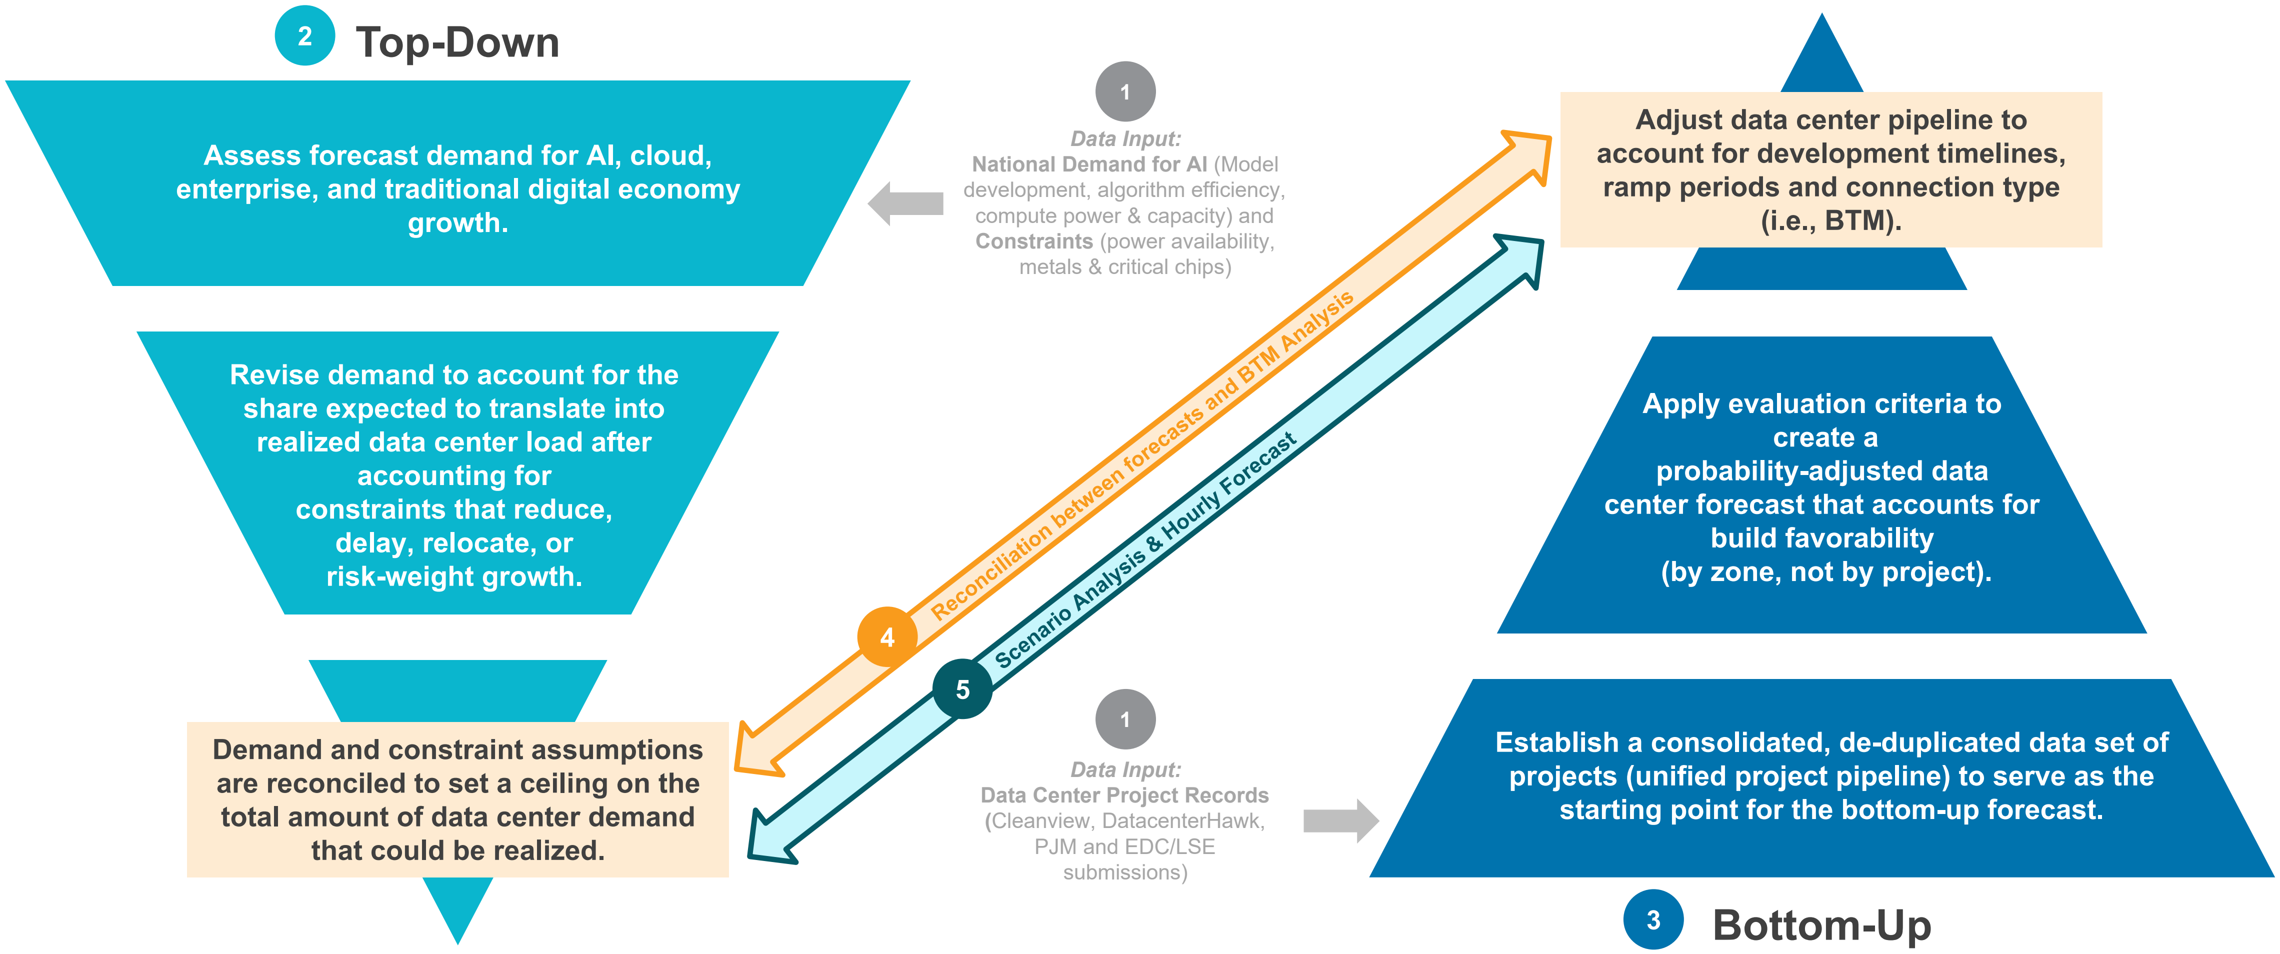

Four layers work together: demand sets the size, the pipeline sets the location, constraints keep it realistic, and behind-the-meter absorbs the overflow

Top-Down

Sets how much capacity is needed in total across the horizon.

Bottom-Up

Sets where it lands, using the actual project pipeline zone by zone + possibility of ignition.

Constraints

Keep the outlook grounded in what the grid can physically absorb.

Behind-the-Meter

Acts as the relief valve when interconnection cannot keep pace.

Together, they turn a national demand number into a credible, **zone-level build path.**

Seven sequential steps reconcile unconstrained demand with zonal and PJM limits, then capture the final build

1

STEP 1

Screen the unconstrained case

Check the unconstrained build for other binding constraints and correct them before allocating to data center resources.

2

STEP 2

Add or ramp projects

Ramp online projects, then advanced, then generic, to meet constrained demand. Relative likelihood is derived from the capture equation on slide 43.

3

STEP 3

Enforce zonal and PJM limits

Both zonal and PJM constraints are applied as projects are added or ramped.

4

STEP 4

Spill beyond PJM

Load still unserved moves to non-PJM zones, subject to available capacity.

5

STEP 5

Shift to behind-the-meter

Where constraints block demand, go BTM to relative fuel scores, within PJM limits.

6

STEP 6

Return BTM to the grid

Bring behind-the-meter load back onto the system after 2 years if capacity exists. Assume island projects return to hybrid service after 2 years.

7

STEP 7

Shape and capture results

Generate hourly shapes including the impact of hybrid, then capture final results.

Behind-the-Meter offers a critical relief for constrained power systems

ALLOCATED ZONAL DATA CENTER LOAD *Tested against the zone's ability to serve it at the 1-in-10 planning standard*

GRID-CONNECTED	HYBRID	ISLAND	SITED ELSEWHERE	DEFERRED
When it happens The zone can serve the load at the 1-in-10 standard **What it means** No behind-the-meter capacity is required **In the forecast** Load appears in full and is served by the grid	**What it is** Grid-connected, self-supplying part of the time **Why chosen** Preferred wherever a connection can be obtained — grid energy is cheaper in most hours, and the grid clears faults and supports restart **Sizing** ~1.11× load. No reserve margin, as the grid provides redundancy **In the forecast** Load remains, served under firm service **Capacity treatment** Credited at its accreditation factor; peak and energy assessed net of behind-the-meter dispatch **Populations** Built-as-hybrid sites take firm service; reconnected islands return on non-firm service and are curtailable	**What it is** No grid service; the site self-supplies its full load **Why chosen** The forced outcome where grid service cannot be obtained on the required timeline — not a preferred configuration **Sizing** ~1.67× load. Redundancy for forced outages, with no external backstop **In the forecast** Neither the load nor its generation appears on the system **Capacity treatment** Not credited, as the resources are absent from the system **Populations** A single population, converting to hybrid on reconnection	**What it is** The project is developed in a different zone **How it arises** Where power cannot be obtained on the required timeline, developers select an alternative site while still in planning **Important** A siting choice made **before** investment; no announced project is moved **Driver** The next-best zone by attractiveness score, PJM first, then beyond	**What it is** The residual — served by none of the other outcomes **What happens** Deferred and re-tested against the following year's supply position **Clears when** New supply, transmission relief or improved build conditions arrive

The split between the two follows observed proportions in the announced project pipeline

RECONNECTION — both behind-the-meter outcomes are temporary: islanded sites seek grid service once the zone can accommodate them, approximately two years, returning on non-firm service and retaining their generation. That capacity re-enters the zonal capacity position as the load returns to grid-served load, lifting net load in the later years of the forecast.

Island and hybrid configurations differ in whether grid connection is retained, determining how each is sized and counted

HYBRID

Grid-connected — self-supplies part of the time

Sizing

Power load served = 1.00×

generation installed = 1.11× (ICAP)
= 0.69x (UCAP)

No reserve margin – the grid provides redundancy.

Why this configuration

Preferred wherever a connection can be obtained. Grid energy is cheaper in most hours, the option to purchase retains value even at low utilization, and the grid clears faults and supports restart.

In the load forecast

The load remains and is served under firm service; the generation offsets part of it.

In the capacity position

Credited at its accreditation factor, with peak and energy assessed net of behind-the-meter dispatch. Credited because sites are built at approximately one-to-one with no reserve margin, so they cannot be assumed reliably interruptible — and because the technology mix points to a capacity rather than an energy motivation.

Populations

Sites built as hybrid take firm service and operate for arbitrage and coincident-peak management. Islanded sites that reconnect return on non-firm service and are curtailable.

ISLAND

No grid service — self-supplies its full load

Sizing

Power load served = 1.00×

generation installed = 1.67 × (ICAP)
= 1.14x (UCAP)

Redundancy for forced outages, with no external backstop.

Why this configuration

The forced outcome where grid service cannot be obtained on the required timeline — not a preferred design. An islanded site cannot lean on the grid to clear faults or restart, which is materially harder at scale.

In the load forecast

Neither the load nor its generation appears on the system until reconnection.

In the capacity position

Not credited, as the resources are absent from the system. The site returns to the forecast on reconnection, which is assumed approximately two years after the zone can accommodate it, on non-firm service and retaining its generation.

Populations

A single population, which converts to hybrid on reconnection.

COMMON TO BOTH: the split between the two follows observed proportions in the announced project pipeline, and all behind-the-meter capacity is temporary – a bridge held only until grid service becomes available.

Fuel scores – how much BTM gas can physically and legally be built?

What the scores are

- The score represents a zone-year assessment of whether gas-fired behind-the-meter generation can physically and legally be built — not whether anyone wants to build there, which is handled separately.
- Three sub-factors, each scored 0–3, averaged into a composite: gas access, build policy, and deliverability. Build policy has double weight because it is the most important.
    1. *Gas access* – distance to interstate mainline with spare capacity, and mainline utilization. Sourced from pipeline GIS and capacity postings
    2. *Build policy* – state and local permitting posture toward on-site generation and associated pipeline infrastructure, including moratoria and litigation history
    3. *Deliverability* – winter basis spread at the serving hub versus Henry Hub, and frequency of operational flow orders or curtailment notices
- A separate binary buildable-policy flag gates the zone entirely; where it is zero, no behind-the-meter generation is permitted regardless of the other scores.

Sample Fuel Scores

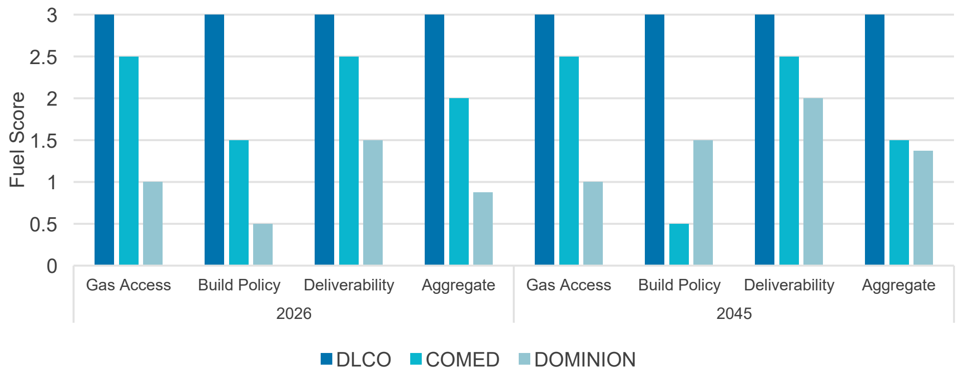

Fuel Scores Rubric

	What it captures	0	1	2	3
Gas Access	Is the gas physically there and can pipe reach the site? (basin proximity + pipeline headroom)	Pipeline island, far from supply	Lateral-dependent, tight	Mainline accessible	On/near basin, mainline with spare
Build Policy	Can you permit and build the pipe/lateral the gas facility?	Moratorium / de facto ban	Hostile, litigated	Mixed/slow	Supportive
Deliverability	Will gas show up when needed? (firm availability + winter constraint risk)	Regular curtailment	Tight, interruptible-only	Workable	Firm available, low winter risk
Buildable Policy	Can BTM generation be built at all?	No	Yes	N/A	N/A

CRA Charles River Associates

Projecting the Behind-the-Meter build limit

As with grid-connected resources, the ability to build BTM resources is assumed to grow over time

Gas Build Limit

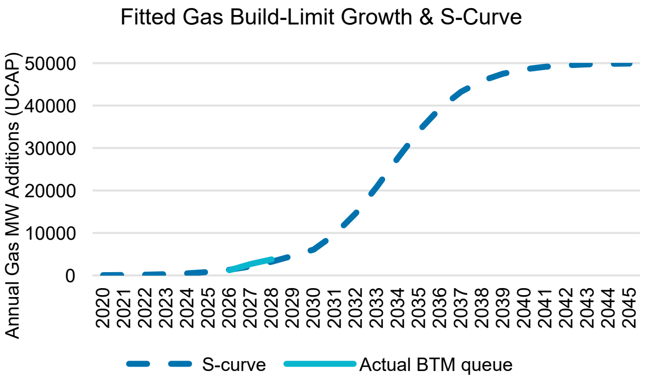

Storage Build Limit

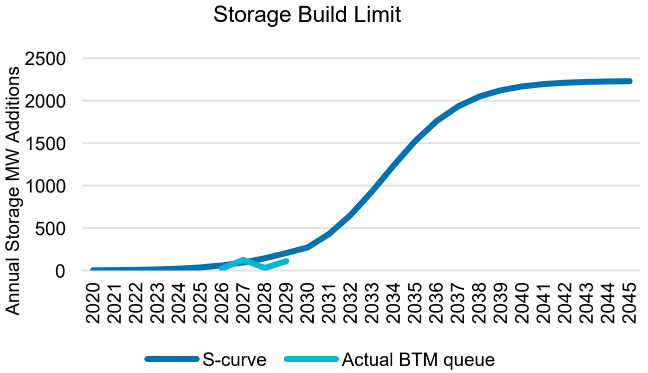

Note: calculated as the change in observed MW capacity between each month and the same data center's capacity 6 months earlier, averaged across the ramp period (which is between the first MWs coming online, and when the data center's MWs are equal to the maximum observed in the data)

Approach

Applied PJM-wide, not zonally

Announced BTM projects with a commercial online date, aggregated by year and used as an annual flow

Absent a formal queue, the pipeline is the equivalent revealed-preference signal

Joined to US turbine orders under 100 MW, 2020–25 — same manufacturing base, so the longer history informs the shape

The junction carries the last observed growth rate, so the splice adds no step; the fitted logistic levels off rather than compounding

Source: LSE Submissions

Assumed BTM ramp rate: BTM projects ramp slower than grid-connected projects

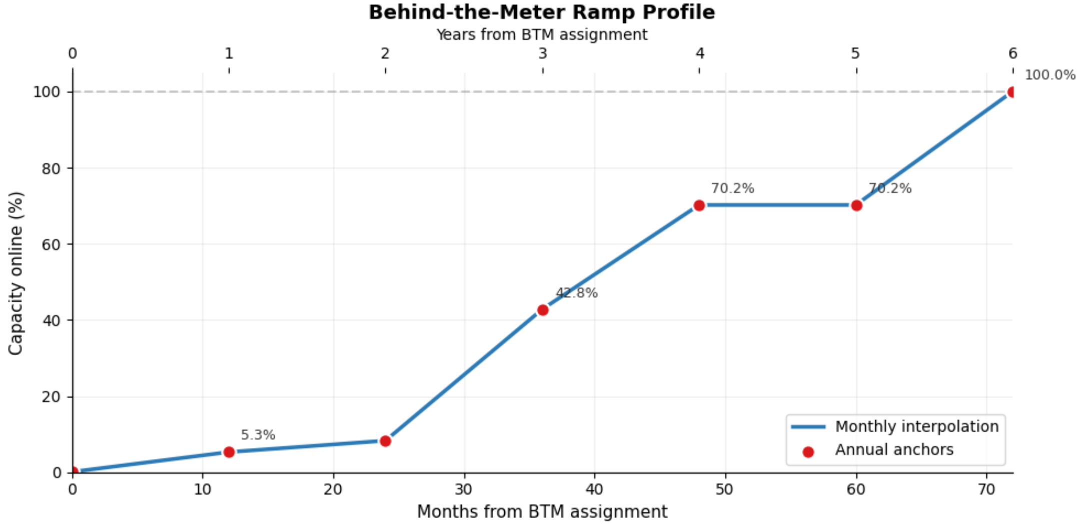

Example where all data center campuses start their ramp at the same time

Assumptions

- Assumes ramp rates follow year-over-year build trends as observed in public announcements
- Build rates are slower than grid-connected projects due to the need to self-build power

Source: CRA analysis, Cleanview

Data centers with BTM generation will ultimately seek grid reconnection

Key Assumptions

1. Islanded behind-the-meter build allows data centers to come online *early*, but it is *expensive* relative to grid service.

2. Data centers without utility connection are assumed to seek grid service as *soon* as it becomes available, consistent with a connection being preferred wherever it can be obtained.

3. Islanded sites cannot be added to hybrid demand in any year that zone is *constrained*, as proxied by CETO/CETL limits in each year.

4. Islanded sites are assumed to try to become hybrid *seven years* after demand is forecasted, subject to above constraint. This is the average time until a major overhaul would be required on the existing generating resources.
    - BYOG connection timelines are assumed to be shorter than full-service data centers (which average 4-7 years depending on zonal constraints)

5. Sites return on non-firm service, retaining their generation, and are therefore *curtailable* following reconnection.
    - Consistent with observed trends in existing data centers connecting to grid service with existing backup generation

6. Reconnected capacity re-enters the balance and the associated load returns to grid-served load, producing an *upward adjustment* to net load in the later years of the forecast net of on-site generation.

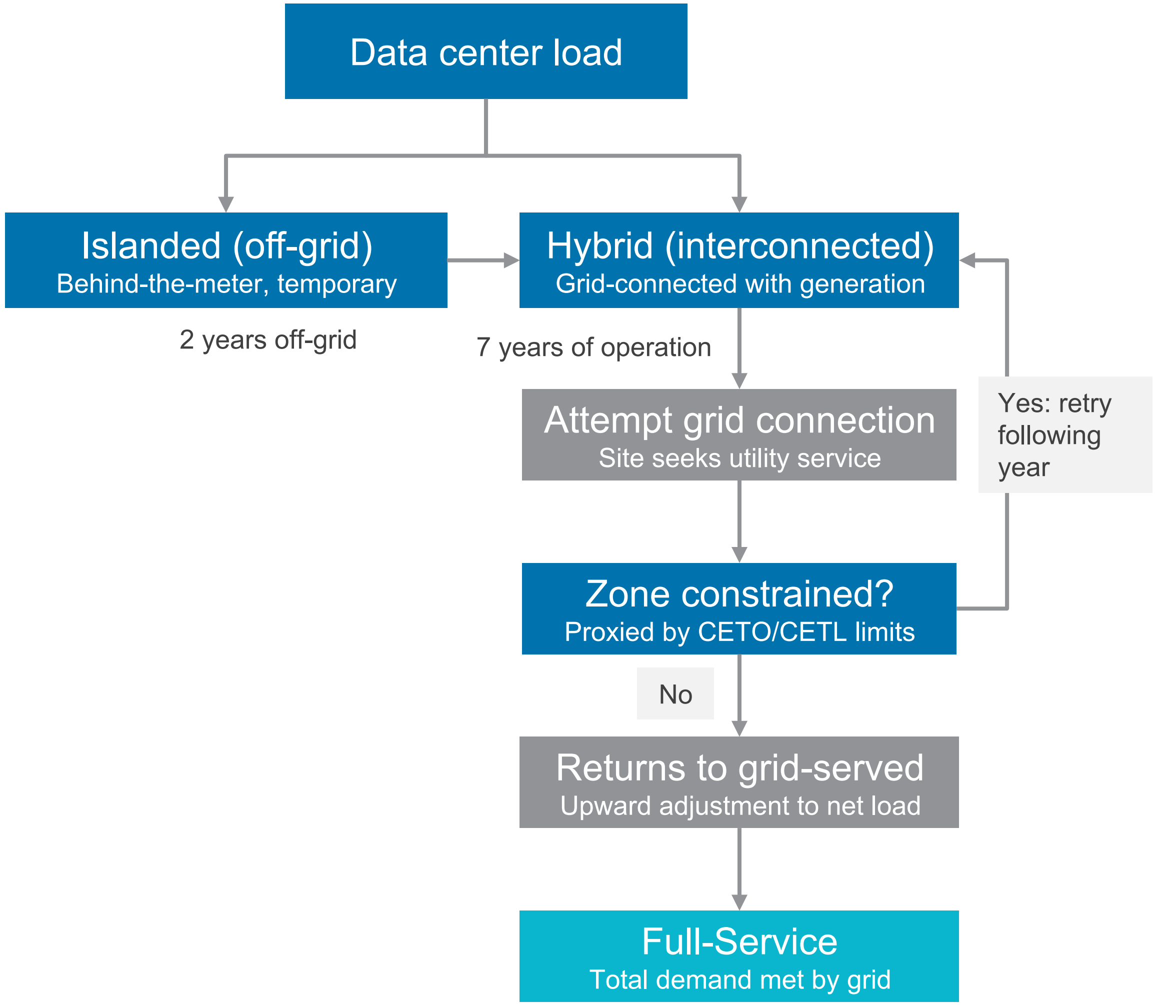

Source: IEA, Lawrence Berkeley

Economic dispatch modeling yields hour-to-hour economic dispatch of BTM resources

1. Grouping tranches
Hybrid BTM resources are grouped into technology-specific tranches by zone. Each tranche is entered into a production cost model that **reduces the gross BTM datacenter load in its zone** — it does not compete as a supply offer in the broader market and carries no capacity accreditation.

2. Dispatch logic
In each hour, model compares each marginal cost to the zonal clearing price. If marginal cost < zonal price, the resource tranche dispatches, reducing datacenter grid draw by that amount. If marginal cost ≥ price, the tranche stays off and the load is served from the grid.

3. Islanded resources
Islanded data centers have no electrical interconnection and are therefore excluded from the zonal demand forecast, reflecting that they would have no impact on net system demand. They only enter the model once they interconnect, at which point they are treated as hybrid resources (see prev. slide).

Key Takeaway:
Because tranche dispatch is price-driven rather than a fixed schedule, the **resulting zonal price series approximates when the operator would turn on its FTM resource**: self-generate when it is cheaper than buying, draw from the grid when it is not.

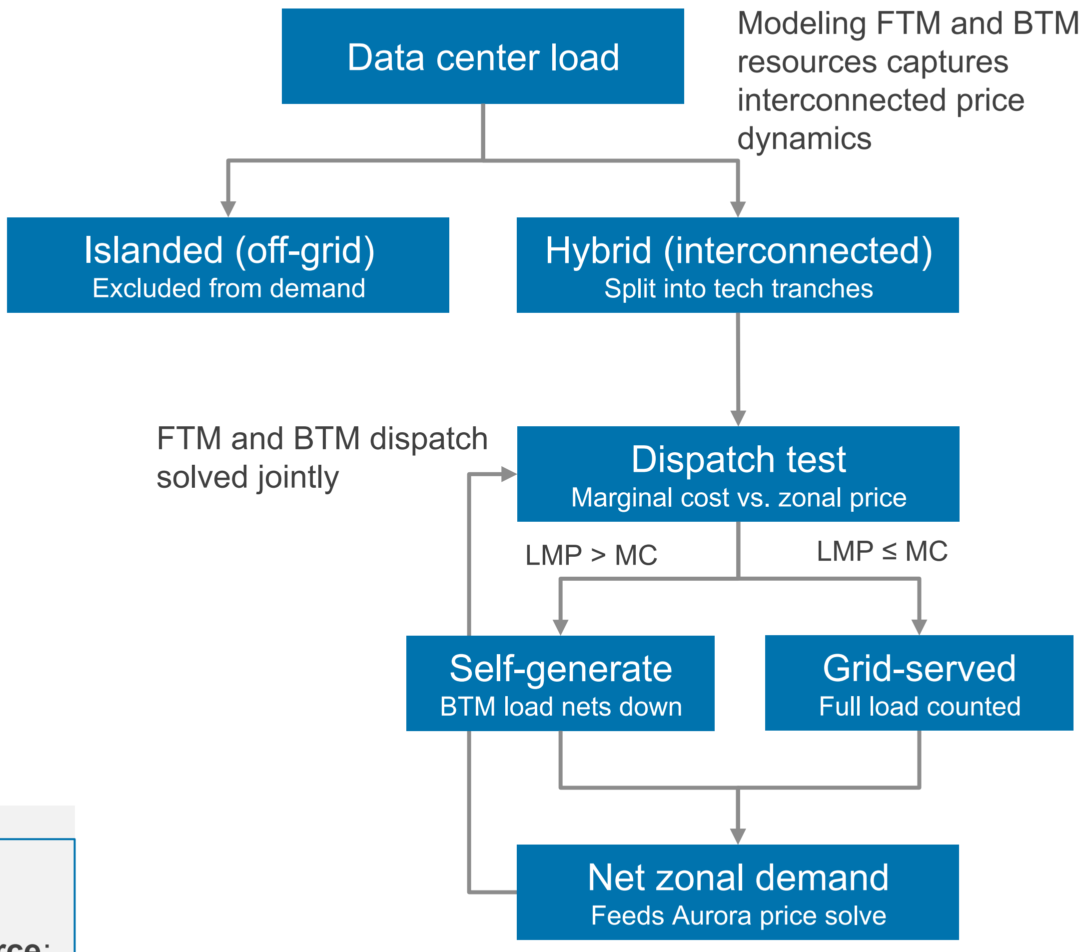

*Zonal price and dispatch are solved together hourly.

Meeting agenda

Topic	Timing
Introductions and July 13 LAS Meeting Recap	9:10 – 9:20 AM
Top-Down: Unconstrained View	9:20 – 9:35 AM
Top-Down: Constraints	9:35 – 10:00 AM
Bottom-Up Forecast	10:00 – 10:15 AM
Load Profile Assumptions	10:15 – 10:25 AM
Forecast Reconciliation and BTM	10:25 – 10:45 AM
Scenarios	**10:45 – 11:00 AM**
Break	*11:00 – 11:10 AM*
Q&A	11:10 – 11:55 AM
Next Steps & Close	11:55 – 12:00 PM

Impact of Scenarios

Using scenarios can illustrate how demand growth and constraints interact under different states of the world, how quickly constraints relax, and therefore how much national demand can realistically translate into regional load

State of the World	Constraint Environment	Impact on Top-Down Forecast
Base Case / Managed Growth	AI adoption continues, but power, supply chain, and siting constraints slow realization. Going forward, new infrastructure accelerates beyond current rates. →	*Demand grows steadily, but constraints cap near-term load realization and push some growth into later years.*
Abundance Case	Infrastructure, compute, and power constraints are relaxed faster than base case expectation; efficiency gains also improve. →	*Faster load growth enabled by fewer binding constraints limiting demand realization.*
Power-Constrained Growth & Evolving Inter-region dynamics	AI demand remains strong, but power availability and interconnection remain severely constrained (current rates remain), particularly in legacy markets. →	*National demand exists, but regional load realization is significantly limited by deliverability, cost, and timing of power. Data centers are likely to increasingly shift toward emerging and tertiary markets.*
Slower Adoption	AI adoption is more measured due to cost, use-case uncertainty, efficiency improvements, or customer adoption friction. →	*Lower overall demand trajectory, with less pressure on infrastructure constraints and slower regional siting.*

Scenarios test a range of internally consistent outcomes

Scenario	S-D Outlook	AI Adoption	AI Efficiency	BTM Generation
Abundance	More generation added than base	Same as base	Same as base	Same availability as base, but less usage
Constrained	Less generation added than base	Same as base	Incentivized to innovate: Increased efficiency gains	Build limits tighten, but larger shortfalls yield more BTM
Low Adoption	Same as base	Lower adoption rate than base case	Same as base	Build limits unchanged, less shortfall so less BTM

Most significant implication

Supply build-out is determined by multiple levers

While near term queue-based outlook remains the same, changes are applied to LTCE forecast (2031 – 2047)

Scenario	Gas build limit	Storage build limit	Cycle 1 treatment	Post-2038 constraint and transmission logic
Abundance *More generation added than base*	Increase to 18% CAGR	Quicker ramp and higher constraint	Decreased attrition represented by build limit relaxation	Off
Base Case *Base case*	Remain at 10% CAGR	Ramp to 6 GW and remain flat	Included as generic units in LTCE	Off
Constrained *Less generation added than base*	Decrease to 5% CAGR	Minimal increase in constraint, increase to ~3.5 GW	Lower levels of additional resources via build limits	Off

Charles River Associates

Turbine and battery supply chains adapt more quickly to increased demand

Supply Forecast Details *for Abundance Case*

- A relaxed turbine constraint is the key lever to achieving this scenario and enabling an accelerated gas build-out; an increased storage constraint allows for additional batteries to come online
- Medium-term gas additions are built sooner due to relaxed constraints, reducing UCAP shortfall by up to ~6 GW annually
- However, near-term shortfalls (2027–2030) remain unchanged as the queue remains a fixed addition

In UCAP	Base Case	Abundance Case
Highest Shortfall	28 GW	28 GW
Cumulative Shortfall (2031–2047)	125 GW	<125 GW
Gas Additions (2031–2047)	53 GW	>53 GW

CRA "Abundance Case" Gas Turbine Constraint (ICAP)

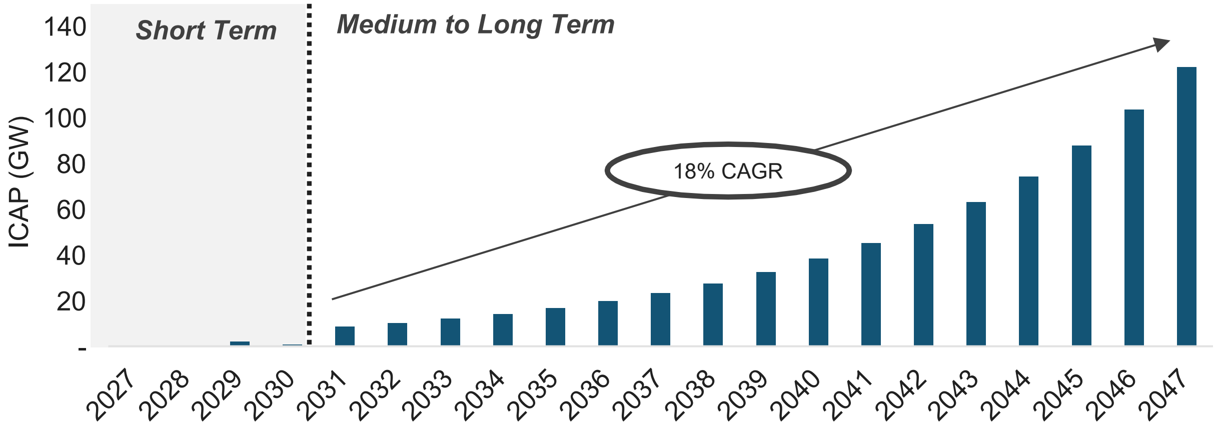

CRA "Abundance Case" Battery Constraint (ICAP)

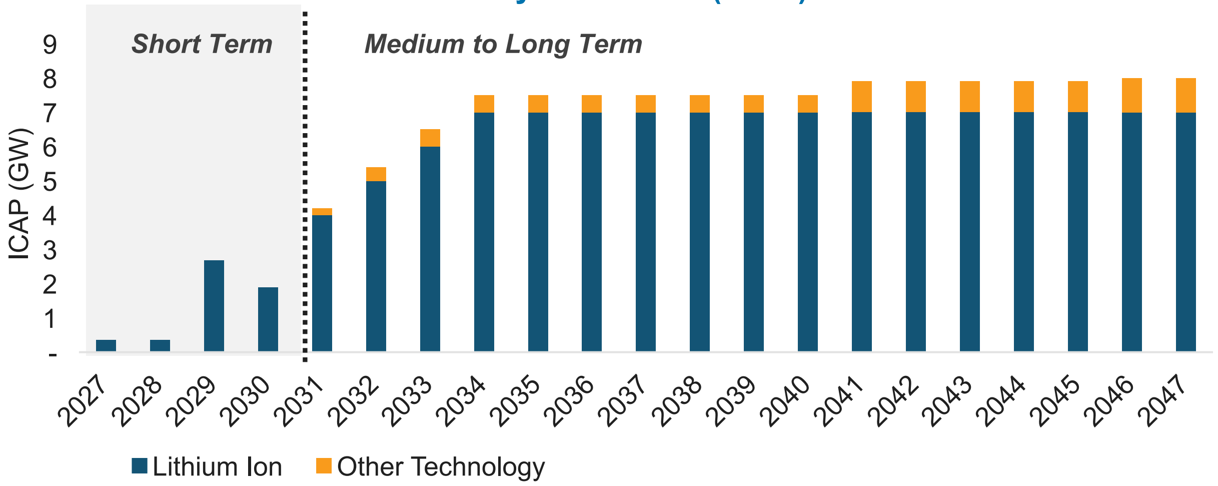

Sources: Utility Dive GE Article; GE Vernova; S&P Global Article; CAISO 20 Year Transmission Outlook

CRA Charles River Associates

Turbine and battery supply chains lag even as demand grows

Supply Forecast Details *for Constrained Case*

- A tightened turbine constraint along with a lower battery constraint prevents new generation from coming online at the same rate as other cases
- Medium-term gas additions are built later due to harsher constraints, increasing UCAP shortfall by up to ~16 GW annually
- However, near-term shortfalls (2027–2030) remain unchanged as the queue remains a fixed addition

In UCAP	Base Case	Constrained Case
Highest Shortfall	28 GW	28 GW
Cumulative Shortfall (2031–2047)	125 GW	>125 GW
Gas Additions (2031–2047)	53 GW	<53 GW

CRA "Constrained Case" Gas Turbine Constraint (ICAP)

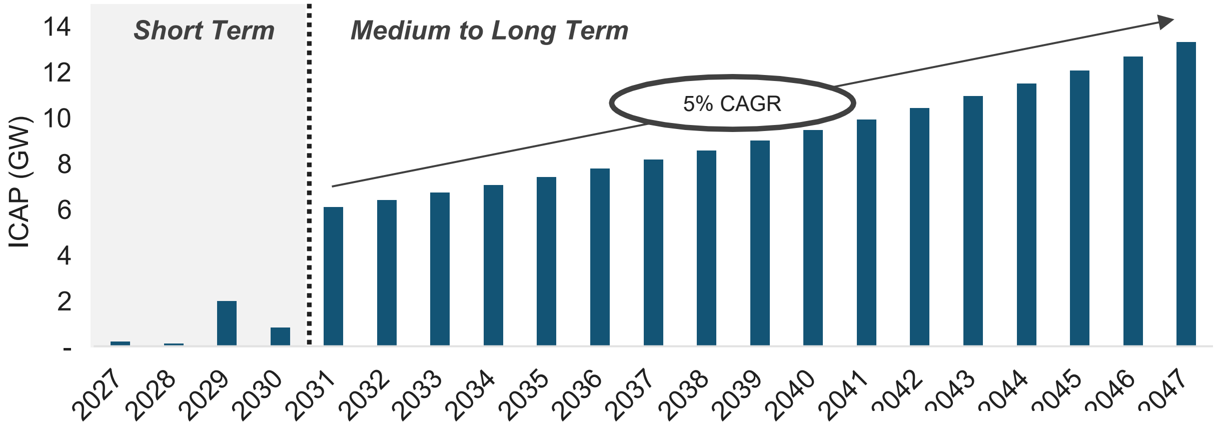

CRA "Constrained Case" Battery Constraint (ICAP)

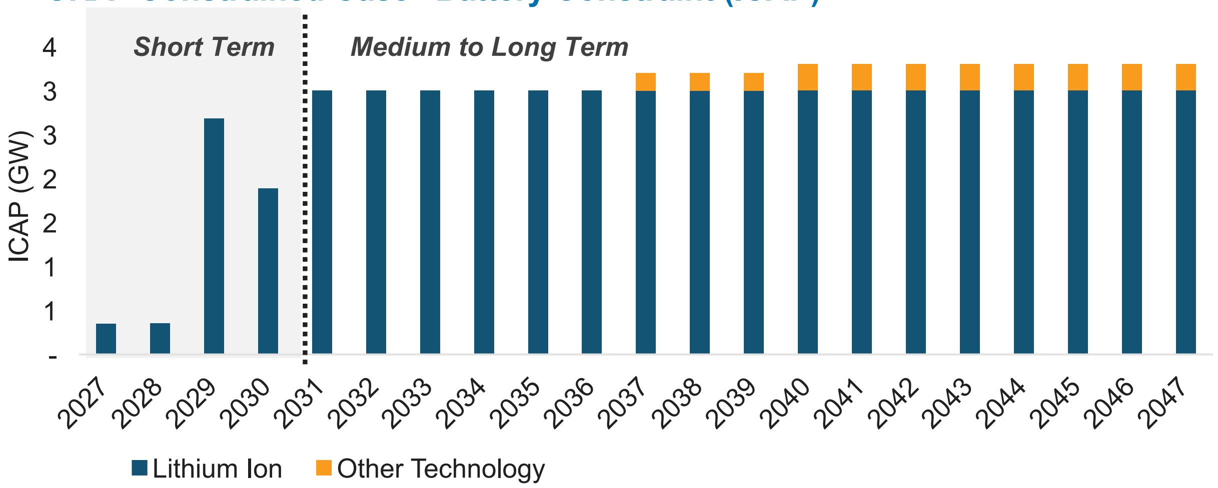

Sources: Utility Dive GE Article; GE Vernova; S&P Global Article; CAISO 20 Year Transmission Outlook

Scenarios test a range of internally consistent outcomes

Scenario	S-D Outlook	AI adoption	AI efficiency	BTM generation
Abundance	More generation added than base	Same as base	Same as base	Same availability as base, but less usage
Constrained	Less generation added than base	Same as base	Incentivized to innovate: Increased efficiency gains	Build limits tighten, but larger shortfalls yield more BTM
Low Adoption	Same as base	Lower adoption rate than base case	Same as base	Build limits unchanged, less shortfall so less BTM

Most significant implication

CRA Charles River Associates

Forced innovation case represents a scenario where AI demand remains strong, but power constraints are persistent

What Needs to be True in the “Forced Innovation” Case

- AI adoption remains strong as companies realize productivity gains from AI
- Power constraints increasingly limit data center expansion and deployment of power-intensive hardware
- Market pressure shifts to upstream suppliers, incentivizing chip manufacturers to focus on computationally efficient designs, increasing chip processing efficiency
- Efficient chip design enables AI adoption while counteracting MW demand growth

	Base Case Growth	Constrained Case Growth
Adoption	95% Individual Adoption by 2040	*No Change*
Chips	1.34x/year	1.39x/year (+5%)

CRA “Forced Innovation Case” Data Center Forecast (Cumulative GWs)

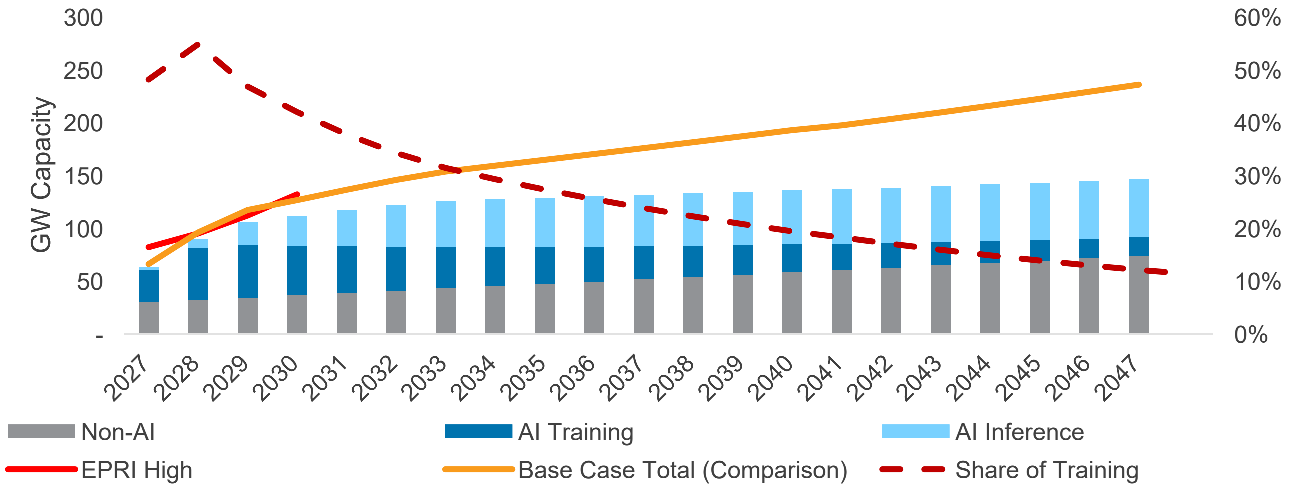

Growth in Base vs Constrained Case Chip Efficiency

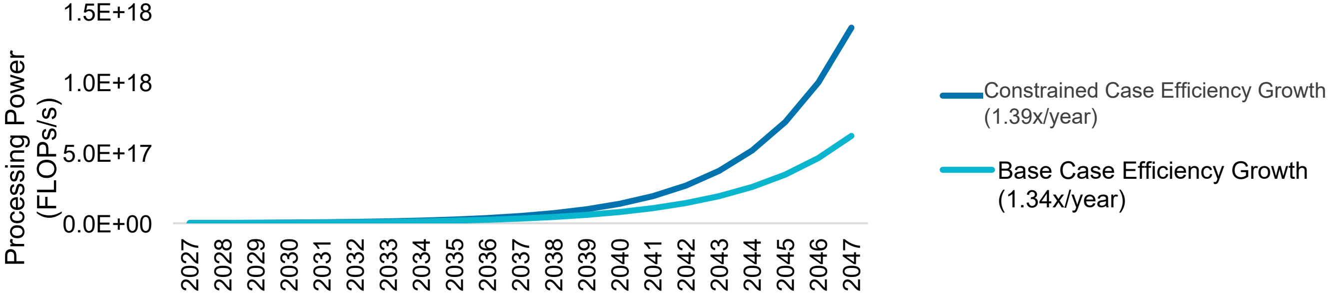

Sources: Epoch.ai

Lower adoption case represents a scenario where AI struggles to gain traction, causing a decline in demand

‘What Needs to be True’ in the Lower Adoption Case

- Enterprise AI adoption slows as productivity gains fall short of expectations.
- Weaker demand reduces AI development investment and model development activity.
- Fewer breakthrough models further dampen adoption, creating a self-reinforcing slowdown in AI growth.

	Base Case Growth	Lower Adoption Case Growth
Adoption	95% Individual Adoption by 2040 80% Enterprise Adoption by 2040 11% CAGR in AI Models Trained	60% (-35%) Individual Adoption by 2040 55% (-25%) Enterprise Adoption by 2040 5.5% CAGR in AI Models Trained (-50%)
Chips	1.34x/year	*No Change*

CRA “Lower Adoption Case” Data Center Forecast (Cumulative GWs)

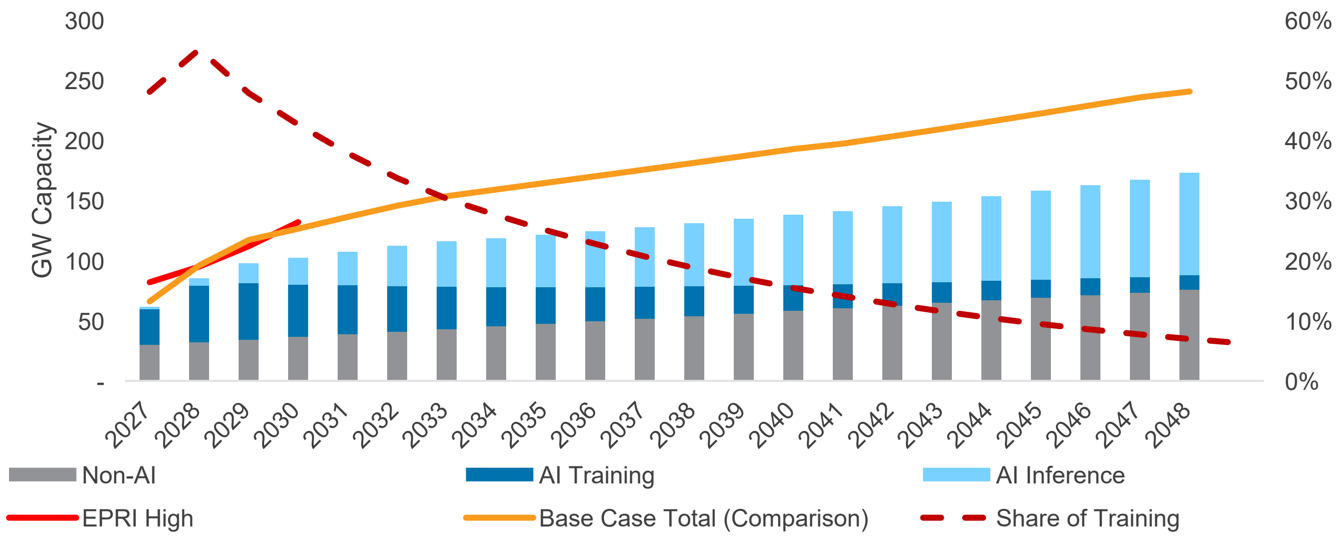

CRA “Lower Adoption Case” Adoption Forecast (% U.S. Population)

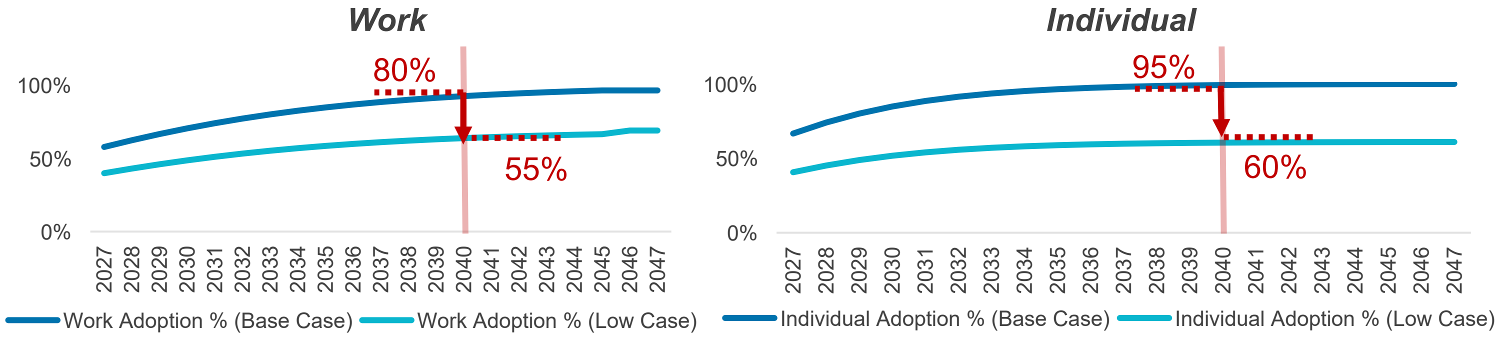

Lower Adoption Case Adoption Curve is based on JoA study suggesting 31% of individuals do not see clear benefits from using AI at work. The constrained case curve is capped by the % of individuals who did not report benefits from using AI at work. Gallup Poll indicates that 39% of individuals see AI doing more harm than good. Individual usage is capped by the % of individuals that see AI doing more harm than good.

Meeting agenda

Topic	Timing
Introductions and July 13 LAS Meeting Recap	9:10 – 9:20 AM
Top-Down: Unconstrained View	9:20 – 9:35 AM
Top-Down: Constraints	9:35 – 10:00 AM
Bottom-Up Forecast	10:00 – 10:15 AM
Load Profile Assumptions	10:15 – 10:25 AM
Forecast Reconciliation and BTM	10:25 – 10:45 AM
Scenarios	10:45 – 11:00 AM
Break	***11:00 – 11:10 AM***
Q&A	11:10 – 11:55 AM
Next Steps & Close	11:55 – 12:00 PM

Meeting agenda

Topic	Timing
Introductions and July 13 LAS Meeting Recap	9:10 – 9:20 AM
Top-Down: Unconstrained View	9:20 – 9:35 AM
Top-Down: Constraints	9:35 – 10:00 AM
Bottom-Up Forecast	10:00 – 10:15 AM
Load Profile Assumptions	10:15 – 10:25 AM
Forecast Reconciliation and BTM	10:25 – 10:45 AM
Scenarios	10:45 – 11:00 AM
Break	*11:00 – 11:10 AM*
Q&A	**11:10 – 11:55 AM**
Next Steps & Close	11:55 – 12:00 PM

Meeting agenda

Topic	Timing
Introductions and July 13 LAS Meeting Recap	9:10 – 9:20 AM
Top-Down: Unconstrained View	9:20 – 9:35 AM
Top-Down: Constraints	9:35 – 10:00 AM
Bottom-Up Forecast	10:00 – 10:15 AM
Load Profile Assumptions	10:15 – 10:25 AM
Forecast Reconciliation and BTM	10:25 – 10:45 AM
Scenarios	10:45 – 11:00 AM
Break	*11:00 – 11:10 AM*
Q&A	11:10 – 11:55 AM
Next Steps & Close	**11:55 – 12:00 PM**

Next Steps

- EDC/LSE form submissions due – September 4
- Individual meetings with EDCs/LSEs in September/October
- Final results presentation October 14 LAS meeting

We thank you for your support and participation, both today and throughout this process!

For any additional questions or comments please reach out to Load_Analysis_Team@pjm.com.

CRA Charles River Associates

CRA Charles River Associates