MARKET BRIEF
JULY 2026

GRID RELIABILITY & ENERGY FLEXIBILITY

When the Grid Is Tightest

The 2026 Northeast heat dome, market mechanics, and the value of operated flexibility

Abstract. The heat dome settling over eastern North America in late June and early July 2026 is not merely a weather event; it is a synchronized reliability stress test arriving on a grid whose reserve margins have thinned structurally. Record or near-record demand is landing simultaneously across four of the continent's largest markets at a moment when data-center load, electrification, and generator retirements have outpaced new supply. For large energy users, the event is best understood not through the lens of consumption, but through the cost-allocation mechanisms it activates — capacity charges, Global Adjustment, ICAP tags, and scarcity pricing — each of which is set by demand in precisely the hours now under stress. This brief reviews the market mechanics at work, positions behind-the-meter battery storage as the flexibility instrument best suited to heavy industry, and sets out how an operating partner converts that exposure into managed cost and grid-reliability value.

THE STAKES, IN THREE NUMBERS

~$120K
annualized capacity cost per MW-year in PJM at the current auction cap ($329/MW-day × 365)

~$335K
indicative annual Global Adjustment cost per MW of coincident-peak demand for an Ontario Class A facility

1 hour
the single NYCA peak hour that sets a New York meter's capacity tag for the entire capability year

Derivations and per-market detail in Section 3.4. Figures are indicative; see table notes. Prepared by the Rodan Energy Markets team.

1 A synchronized reliability event

A persistent heat dome has placed more than 250 million people under extreme-heat conditions across the Plains, the Midwest, and the eastern seaboard. The operative feature for grid operators is not the peak temperature on any single afternoon but the *persistence and geographic breadth* of the event: sustained heat drives cooling load higher and holds it there, and because overnight temperatures remain elevated, facilities begin each day from a warmer baseline, forcing cooling systems to start earlier and run longer.

What distinguishes this event is that the stress is *coincident across markets*. When demand peaks in the same late-afternoon and early-evening windows across PJM, New York, MISO, and Ontario at once, the regional diversity that normally lets neighboring systems lean on one another erodes precisely when it is most needed. MISO illustrates the systemic point most cleanly: the footprint's geographic breadth is normally its strength, and a continental heat dome removes it. When zones peak simultaneously, internal diversity disappears, reserves tighten across the interconnection, and real-time energy and capacity-related costs rise for a broad industrial base in the same moment.

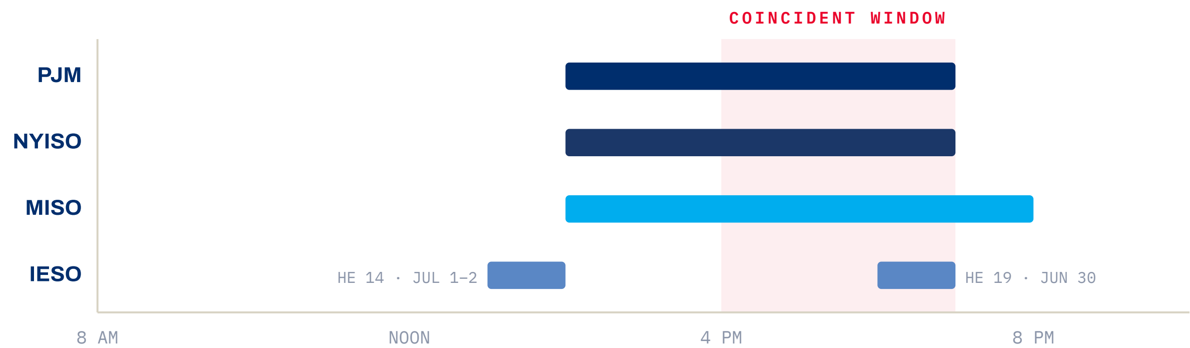

Figure 1. Typical peak windows during the event, by market. When demand peaks in the same late-day hours everywhere at once, regional diversity — the grid's normal shock absorber — disappears. Windows are indicative, drawn from operator advisories and the Rodan Energy Markets team's coincident-peak tracking (pending ISO verification).

MARKET	RELIABILITY SIGNAL THIS WEEK	HOW IT CONVERTS TO BUYER COST
PJM	RTO demand peaked at ~161,859 MW on July 1 (HE 18) and ~157,031 MW on July 2 (HE 15)†; EEA-1, Hot Weather / Max-Gen / Load-Management alerts; DOE 202(c) emergency orders.	Capacity cleared near the ~$329–333/MW-day cap for a third straight auction, ~6,625 MW below the 20% reserve-margin target; scarcity feeds real-time LMP and 5CP charges.
IESO	New season peak of 23,833 MW on July 1 (HE 14), ahead of 23,526 MW on June 30 (HE 19) and 23,438 MW on July 2 (HE 14) — three top-five candidates in three days†.	Class A Global Adjustment is set by Peak Demand Factor — a facility's share of the top-five coincident provincial peaks — allocating a ~$7.9B (2024) and rising GA pool for 12 months.
NYISO	Coincident peak of ~29,582 MW on July 1 (HE 19)†, against a summer margin NYISO called the lowest in recent history (~417 MW; 24.5% IRM).	The single highest NYCA hour in July/August sets each meter's ICAP tag for the following capability year; Zone J scarcity compounds the charge.
MISO	Broad, multi-zone heat that erases the benefit of regional load diversity across a large footprint.	Simultaneous zonal peaks tighten reserves and lift real-time energy and capacity-related costs across a wide industrial base.

Table 1. The June–July 2026 heat dome as a cost event across four markets. †Peak values reflect the Rodan Energy Markets team's coincident-peak tracking as of July 2 and have not yet been verified by the ISO of each jurisdiction. The deep-dives in Section 3 focus on the three markets where coincident-peak cost allocation is most surgical: PJM, Ontario, and New York.

2 Why this event lands on a tighter grid

Record demand is not new. What has changed is the margin between demand and available supply. Three structural forces have compressed it:

Load growth driven by data centers. PJM's projected annualized summer-peak growth has risen to roughly 3.6% per year over the coming decade — more than ten times the 0.3% per year it forecast as recently as 2021 — with the increase attributed almost entirely to large data-center loads. Ontario's system planner projects net summer peak climbing from about 24 GW today toward 36 GW by 2050, driven by electrification, EV manufacturing, and hyperscale computing.

Supply additions lagging retirements. PJM's most recent capacity auction cleared roughly 6,600 MW below its installed-reserve-margin target, with almost no new generation clearing; NYISO entered summer 2026 on a historically thin reliability margin. In both markets, new resources are not arriving fast enough to offset planned retirements and rising demand.

Reliability agencies flagging the risk. The North American Electric Reliability Corporation has issued elevated alerts tied specifically to the unpredictability of large loads such as data centers — a shift from prior years, when comparable alerts centered on the variability of wind and solar output.

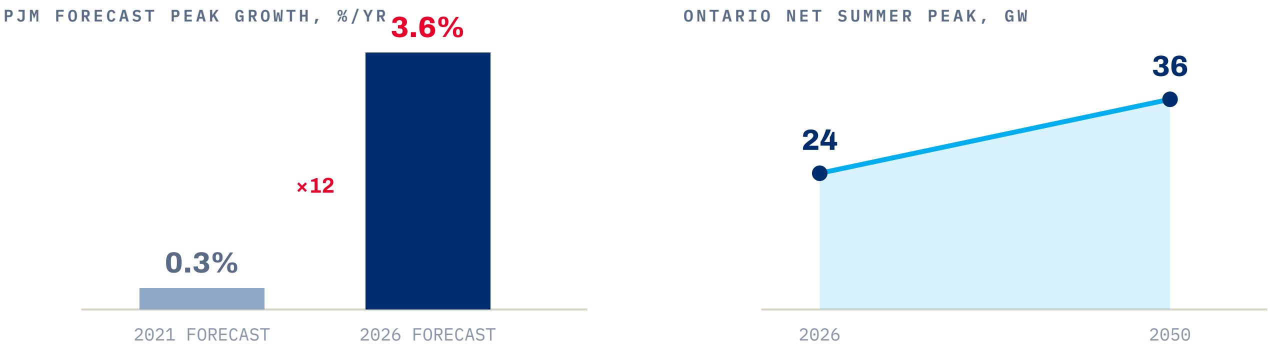

Figure 2. Demand growth has re-rated. Left: PJM's forecast annualized summer-peak growth, 2021 forecast vs. current — a twelvefold increase attributed almost entirely to data centers. Right: Ontario's planned net summer peak trajectory toward 2050.

> ***The same peak megawatt now carries more reliability weight, and more cost, than it did a decade ago. Scarcity that once appeared only in extreme tail conditions is now a recurring feature of hot-weather operations — and it is being priced accordingly.***

3 What the energy markets team sees: from heat to cost

For a sophisticated buyer, the useful question is not "how hot is it," but "which of my forward cost lines is being set this week, and by which hours." Each market answers that differently.

For the finance office, it also helps to name the category correctly. A capacity tag, a Peak Demand Factor, or a 5CP allocation is not a utilities expense that varies with consumption; it is a *twelve-month liability repriced in a handful of hours* with no mechanism to reopen it once set. Most industrial CFOs run disciplined hedging programs for currency, commodity inputs, and interest-rate exposure. Coincident-peak exposure is frequently the largest line in the energy budget that carries no hedge at all, and unlike a financial hedge, the only instrument that manages it is physical: verified load reduction or battery discharge delivered in the setting hours. The sections that follow identify, market by market, which liability is being repriced this week — and Section 3.4 puts a per-megawatt price on leaving it unmanaged.

3.1 PJM — capacity scarcity, made visible

PJM's capacity construct is designed to reveal system tightness through price. Its last auction cleared at, or effectively at, the administrative cap of roughly $329–333/MW-day — a record for the third consecutive auction — while still falling short of the reserve margin needed to hold loss-of-load expectation to the one-event-in-ten-years standard. A negotiated price collar is currently suppressing what would otherwise be a materially higher clearing price. This week's alerts, the EEA-1 declaration, and the federal 202(c) emergency orders authorizing large-load curtailment and pollution-limit waivers are the operational expression of exactly the scarcity the capacity price has been signaling. For buyers, that scarcity flows through in two ways: elevated real-time locational marginal prices during the peak windows, and — in markets and utilities that use coincident-peak accounting — transmission and capacity charges set by demand during a handful of critical hours (the 5CP construct), several of which are being set this week.

Figure 3. Scarcity, priced. PJM's latest auction procured ~6,625 MW below its 20% reserve-margin target while clearing at the administrative price cap for the third consecutive time — with a negotiated collar suppressing a materially higher price.

3.2 Ontario — Global Adjustment and the five hours that set a year

Ontario offers the clearest illustration of why a single hour matters. Class A customers under the Industrial Conservation Initiative pay their share of Global Adjustment — the charge that recovers the cost of contracted generation, refurbishment, and conservation programs — in proportion to their Peak Demand Factor: the facility's contribution to the top five coincident provincial peak hours across the May 1–April 30 base period. That factor then governs the customer's GA allocation for the full following adjustment period, July 1 through June 30.

The consequence is asymmetric. System-wide GA reached approximately $7.9 billion in 2024 and is forecast to rise toward $9.4 billion by 2030 under base-case assumptions; a Class A facility's share of that pool for an entire year is fixed by its metered demand during just five hours. The heat-dome peaks now printing on Rodan's coincident-peak tracker — 23,833 MW on July 1 (HE 14), 23,526 MW on June 30 (HE 19), and 23,438 MW on July 2 (HE 14), pending IESO verification — are strong candidates to become base-period peaks. A facility that fails to curtail through them can carry an elevated GA allocation for twelve months; one that reduces load precisely in those hours locks in a lower factor. Compounding the stakes, the roughly 14% GA relief Class A customers have enjoyed since 2021 is phasing out beginning in 2026, and the introduction of Ontario zonal pricing means commodity cost now varies by hour and by location across all 8,760 hours — extending the optimization problem well beyond the traditional five-peak game.

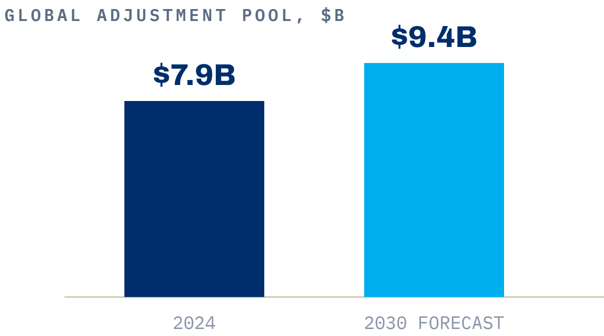

Figure 4. A rising multi-billion-dollar pool, allocated to Class A facilities by their demand in just five hours out of 8,760.

3.3 NYISO — one hour, one tag, one year

New York sharpens the point further. Unlike markets that average several coincident peaks, NYISO sets a meter's installed-capacity (ICAP) obligation from its demand during the *single highest* NYCA system hour in July or August. That one hour establishes the capacity tag that drives capacity charges for the capability year beginning the following May. With the statewide installed reserve margin set at 24.5% against a summer headroom NYISO itself described as the lowest in recent history, and with transmission-constrained Zone J facing peaker retirements that lift in-city capacity prices, the cost of being exposed during the wrong hour this week is measured in a full year of elevated charges.

Figure 5. One hour, one tag, one year. A New York meter's ICAP obligation is set by a single system hour — there is no averaging and no reopening it.

3.4 The cost of one missed hour, per megawatt

The mechanics above reduce to a number a finance team can act on: what one megawatt of unmanaged coincident-peak demand costs per year, and — equivalently — what one megawatt of verified reduction in the setting hours is worth. The figures below are derived from the sources cited in this brief and are intended as planning-grade indications; actual values vary by zone, utility, and rate class.

PJM	**Capacity obligation + 5CP transmission** *Five coincident system peaks, typically summer afternoons*	**≈ $120K** /MW-YR *capacity alone, at the cap — before 5CP transmission, which can be comparable*
IESO	**Class A GA via Peak Demand Factor** *Top five provincial peak hours, May 1 – April 30 base period*	**≈ $335K** /MW-YR *rising with the pool toward 2030; Class A relief phasing out from 2026*
NYISO	**Installed-capacity (ICAP) tag per meter** *The single highest NYCA hour, July or August*	**$100–200K** /MW-YR *at recent Zone J spot capacity prices; lower upstate, pressured upward by peaker retirements*

Table 2. Indicative per-MW annual value of coincident-peak exposure, by market. Bars are proportional to the indicative midpoint. Rodan's energy markets team maintains current per-MW values by zone, utility, and rate class.

4 From exposure to strategy and why execution is the hard part

The most complex energy users in the northeast already understand these mechanisms. Knowing that a peak sets a cost is not the differentiator; acting on it, correctly, in real time is. Capturing the value requires forecasting which hours will set the peaks, deciding what load can move without disrupting production, dispatching the response within the operator's notification window, and reconciling interval data afterward to confirm the reduction landed where it mattered. Each of these is an operational discipline, not a spreadsheet exercise, and each carries execution risk. A curtailment missed by one hour, or a battery dispatched against the wrong signal, forfeits a year of savings.

FORECAST *the setting hours* → **DECIDE** *what load can move* → **DISPATCH** *inside the window* → **RECONCILE** *against published peaks*

MISS BY ONE HOUR → FORFEIT A YEAR

Figure 6. Peak capture as an operating discipline. Every step carries execution risk; the dispatch decision is where a year of savings is won or lost.

This is the gap between flexibility on paper and flexibility that shows up when the grid needs it — and it is where an operating partner earns its place. It also raises the question any finance office should put to a provider in this space: when the dispatch misses the hour, who bears the cost? Section 6 answers it.

5 Positioning battery storage for large industrial

For heavy industry, the limitation of load-based demand response is structural: a production line cannot always be curtailed. Throughput commitments, product quality, process continuity, and customer contracts make shutting a line during a peak either impossible or more expensive than the charge it would avoid. This is why demand response, on its own, underserves the largest and most process-intensive facilities — the very users with the most capacity and GA exposure.

Behind-the-meter battery storage resolves that constraint by *decoupling flexibility from operations*. The battery discharges into the coincident peak while the plant continues running at full rate. Nothing is curtailed; the facility simply draws less from the grid in the hour that sets its cost. For a process load that cannot flex, this is often the only way to participate in peak management at all.

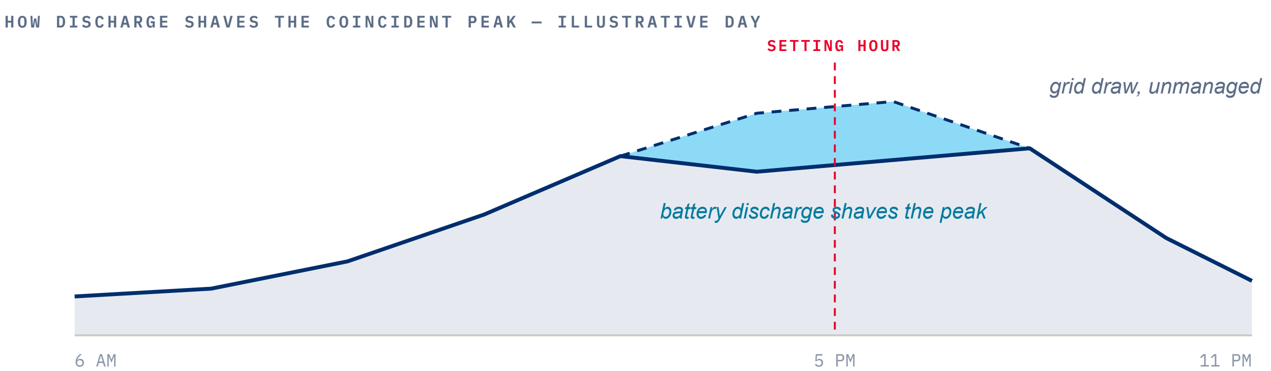

Figure 7. Decoupling flexibility from operations. The plant runs at full rate; the battery discharges through the setting hours, and the meter — not the production line — is what the market sees. Illustrative profile.

The economics improve further because a single asset can stack multiple value streams — the discipline that now defines storage project returns:

Figure 8. Value stacking on a single asset — illustrative composition. The capacity-reduction stream is typically dominant; the remainder is captured only when the asset is actively operated.

Capacity-tag and demand-charge reduction

Energy arbitrage

Discharging through the peaks that set a PJM capacity obligation, a NYISO ICAP tag, or an Ontario Peak Demand Factor directly lowers the single largest, most persistent cost line — one that, once set, holds for a full year (Table 2 puts a per-MW value on it).

Charging in low-price hours and discharging into high-price intervals captures a spread that widens as Ontario's forecast HOEP rises from about $62/MWh in 2026 toward $90/MWh by 2030, and as scarcity lifts real-time LMPs during events like this one.

Wholesale and grid-service revenue

The same asset can earn demand-response payments, ancillary-service revenue, and capacity value by responding when the operator calls — converting an avoided-cost asset into a revenue-generating one.

Resilience

During emergency operating conditions or localized outages — a live risk under this week's EEA and 202(c) posture — stored energy protects critical processes, a value that is real even when no dollars change hands.

The catch is that stacked value is not captured by owning hardware; it is captured by *operating* the asset — forecasting the peak, arbitrating among competing revenue streams hour by hour, and dispatching in real time without compromising the site. A battery run as a passive appliance leaves most of its value unrealized. A battery run as a managed, optimized resource pays for itself several times over. That distinction — hardware versus operated capability — is the entire argument for a partner.

5.1 The finance case: returns, incentives, and structures

The value streams above translate into a returns profile a finance team can underwrite. Three variables dominate it.

Payback. For a well-sited facility with material coincident-peak exposure — an Ontario Class A load, a Zone J meter, a 5CP-billed PJM site — stacked operation of a behind-the-meter battery typically supports paybacks in the 4–7 year range, shorter with incentives, at current capacity and GA price levels, with the capacity-reduction stream alone often covering the majority of the return. Because the charges the asset avoids are set annually and are rising with the market trends described in Section 2, the return profile strengthens rather than decays over the asset's life.

TYPICAL PAYBACK WINDOW, YEARS — STACKED OPERATION

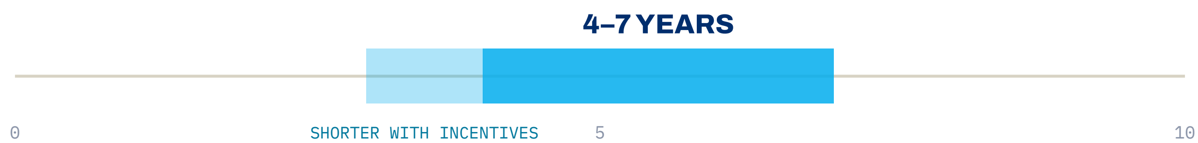

Figure 9. Indicative payback range for a well-sited industrial facility at current capacity and GA price levels; the capacity-reduction stream alone often covers the majority of the return.

Incentives. Jurisdictional incentives materially shorten the payback and should be evaluated site by site. In Canada, the Clean Technology Investment Tax Credit provides a refundable credit of up to 30% of the capital cost of eligible energy-storage property; in the United States, the federal investment tax credit for standalone storage remains available, subject to evolving domestic-content and sourcing rules.

Structure. Ownership is not the only route, and for many industrials it is not the best one. The same asset can be delivered as customer-owned capital, as a financed asset, or under a third-party-capital arrangement in which the operating partner funds the system and is compensated from verified savings and market revenue — no capital outlay from the facility, with performance risk sitting with the party that controls dispatch. Rodan structures and invests in these projects directly: it is a capital-at-risk partner in the outcome, not an advisor recommending someone else's spend. Which structure fits is a treasury decision as much as an energy one — balance-sheet treatment, hurdle rates, and capital-allocation priorities differ — and it is a conversation Rodan is equipped to have with the finance office, not just the plant.

6 Where Rodan fits: Your Energy Team

Rodan develops, optimizes, and operates flexible energy assets across the full chain, from a facility's meter to the wholesale market. The distinction that matters during an event like this one is that Rodan does not advise on flexibility from the sidelines; it operates it in real time, as the client's energy team.

For large energy users. Rodan monitors facility load against live market conditions, maintains demand-response readiness, forecasts the peaks that set capacity, GA, and ICAP costs, and dispatches the response, load reduction where a site can flex, and battery discharge where it cannot. The objective is explicit: protect the facility from demand charges, Global Adjustment and 5CP exposure, and real-time price volatility, without asking production, safety, or quality to give ground.

Who bears the risk of a missed hour. In shared-savings and Rodan-capital arrangements, compensation is tied to verified outcomes, peaks captured, market revenue settled, interval data reconciled against the operator's published peaks, so the party controlling the dispatch decision carries the consequence of missing it. Behind that alignment sits the operating infrastructure that makes the risk manageable:

24×7×365

Real-Time Operations Center staffed through every hour of events like this one

20 years

of coincident-peak seasons refining an algorithmic forecasting and dispatch platform

~1.5 GW

of distributed energy resources under management across the markets in this brief

For utilities and grid operators. Individually, one site's flexibility is marginal. Operated as a portfolio, it becomes a resource the grid can lean on. Rodan aggregates demand response across many industrial and commercial sites, together with operated battery storage, into a dependable, dispatchable reliability resource — one that responds when reserves tighten and load peaks across regions at once, exactly the condition on display this week. This is the aggregation layer that turns thousands of individual assets into grid-scale reliability, delivered by an operator rather than assembled by an intermediary.

> *Rodan is an operator, not an advisor and an operator of aggregated flexibility, not merely an aggregator of it. From a single peak-shaving event to a fully dispatched storage asset, and from one meter to the market, Your Energy Team runs the response.*

7 Before the next event: five questions for your energy and finance teams

Extreme-heat events will keep arriving, and each one sets costs for the year that follows. The following five questions are a working diagnostic. An organization that can answer all five is running flexibility as an operation; in Rodan's experience, most cannot confidently answer more than two.

01 Do we know, in dollars per year, what our current capacity tag, Peak Demand Factor, or 5CP allocation costs us — and exactly which hours set it?

02 Do we have a forecasting-and-dispatch protocol for peak-candidate days, with a defined decision window and a named owner — or do we find out after the operator publishes the peaks?

03 What load can move without touching production, quality, or safety — and where load cannot move, has behind-the-meter storage been evaluated at today's capacity and GA prices rather than the prices of five years ago?

04 If we curtailed or discharged through the wrong hour, how would we know — and who reconciles our interval data against the published coincident peaks?

05 Who, specifically, is accountable for peak capture — and does anyone's compensation depend on it?

8 Conclusion: extreme heat as a planning signal

The hottest days of the year are the most consequential days for grid operations — and, increasingly, for the forward energy budget of every large user. This week's heat dome is testing reliability procedures across four major markets simultaneously, and in doing so it is setting capacity tags, Peak Demand Factors, and coincident-peak charges that will shape costs for the twelve months ahead. As demand growth, electrification, and generator retirements keep reserve margins thin, energy strategy and reliability planning will only converge further. The users who understand their load, manage their peaks, and operate flexible assets — particularly storage — will reduce cost while strengthening the grid they depend on. The others will pay for the same hours, twice: once in the moment, and again for a year.

Turn flexibility into managed cost and reliability value — before the next heat event sets your peaks.

Connect with the Rodan Energy Markets team · rodanenergy.com

REFERENCES & SOURCES

Figures reflect operator advisories, market-operator publications, and public reporting current to early July 2026.

[1] PJM Interconnection — Hot Weather Operations Update and Summer 2026 Outlook (peak-load forecasts; 2006 all-time record of 165,563 MW; ~180,200 MW capacity and ~7,800 MW contracted demand response; EEA and alert postings).

[2] U.S. Department of Energy — Section 202(c) emergency orders authorizing large-load curtailment and environmental-limit waivers for the PJM footprint, effective June 30–July 3, 2026.

[3] PJM — 2026 Long-Term Load Forecast (summer-peak growth ~3.6%/yr vs. 0.3%/yr in the 2021 forecast; data-center-driven demand).

[4] Utility Dive / PJM — 2025 Base Residual Auction results (clearing near the ~$329–333/MW-day cap; procurement ~6,625 MW below the 20% installed-reserve-margin target).

[5] IESO — Global Adjustment, Peak Demand Factor, and Peak Tracker documentation (Class A allocation on top-five coincident peaks; May 1–April 30 base period; July 1–June 30 adjustment period).

[6] Rodan Energy Markets team — coincident-peak tracking, current base period (Ontario demand of 23,833 MW on July 1, 23,526 MW on June 30, and 23,438 MW on July 2, 2026; PJM RTO ~161,859 MW and NYISO ~29,582 MW on July 1, 2026). Values pending verification by the ISO of each jurisdiction.

[7] Industry analysis of Ontario GA trajectory (~$7.9B system-wide in 2024; ~$9.4B base-case by 2030), zonal-pricing implications, HOEP forecast (~$62/MWh 2026 to ~$90/MWh 2030), and the phase-out of Class A relief beginning 2026.

[8] NYISO / NYSRC — Summer 2026 reliability outlook and 2026–27 IRM Study (24.5% installed reserve margin; summer margin described as the lowest in recent history; single-hour ICAP-tag methodology; Zone J constraints).

[9] NERC — elevated reliability alert tied to large-load (data-center) uncertainty, 2026.

[10] Behind-the-meter storage literature (NYSERDA technical assessment; value-stacking analyses) on demand-charge reduction, energy arbitrage, grid-service revenue, and resilience as combinable value streams.

[11] Government of Canada — Clean Technology Investment Tax Credit program guidance (refundable credit of up to 30% for eligible energy-storage property); U.S. Department of the Treasury / IRS — investment tax credit provisions applicable to standalone energy storage.