Exploring PJM's Paths with Long-Term Hedges

PREPARED BY

Samuel A. Newell
Andrew Levitt
Joe DeLosa III

PREPARED FOR

Google

AUGUST 31, 2026

Brattle

NOTICE

- This white paper reflects the perspectives and opinions of the authors and does not necessarily reflect those of Google, Alphabet Inc., or any of the affiliates, nor of the Brattle Group's other clients or other consultants.
- Where permission has been granted to publish excerpts of this white paper for any reason, the publication of the excerpted material must include a citation to the complete white paper, including page references.
- The authors would like to acknowledge valuable contributions from colleagues, including Johannes Pfeifenberger, Kathleen Spees, Nathan Felmus, and James Herrigel.

TABLE OF CONTENTS

I. Executive Summary

In this paper we respond to PJM's "Powering Reliability through Market Design" paper by elaborating on and evaluating the hedging constructs PJM proposed in "Path A" and "Path C."[1] Those Paths aim to solve the "credibility trap" PJM described, whereby merchant investors' anticipation of continued interference in prices frustrates the investment that many states rely on to maintain resource adequacy. We discuss Path B only briefly as a complement to Paths A and C, since it alone is not comparable in scope and does not include hedging needed to reduce the customer spot price exposure at the root of the "credibility trap" PJM identified.

PJM's proposed "Path A" adds mandatory long-term capacity hedging to the existing market construct. "Path C" similarly adds mandatory long-term hedging, but for exposure to high energy prices as the expression of scarcity value instead of capacity. This feature of Path C requires enabling changes to the energy market: more robust energy scarcity pricing through a higher cap on energy prices as part of an enhanced operating reserve demand curve (ORDC), and a Swedish-type strategic reserve mechanism instead of a capacity market to ensure reliability targets are met. Under both Path A and Path C, the long-term hedging would serve the dual purposes of incrementally supporting new resource investments and protecting customers against periodic price spikes.

For Path A and Path C, this paper focuses on the **types and scopes of hedging requirements** and corresponding procurement process that might be most acceptable to LSEs and their customers (including large loads), state agencies, and generation owners (as shown in Figure 1 below). For Path C, we expand on necessary **complementary market design elements** that differ from Path A and the status quo, including changes to the underlying spot market construct, interconnection, and resource adequacy assurance mechanisms. We also outline a possible **transition plan**.

[1] PJM, Powering Reliability Through Market Design: Addressing Rising Demand and Constrained Supply, and Stimulating Investment to Support Durable Reliability (May 6, 2026).

FIGURE 1: HEDGING PATHS AT A GLANCE

	Path A: Capacity Hedging	Path C: Energy Hedging
Complementary market construct	**Existing energy + capacity market** Status quo construct	**Energy-only market** Enhanced ORDC with high price cap
RA assurance	**Existing capacity market** RA requirement with capacity market providing "missing money"	**Strategic reserve backstop*** Residual RA needs met out-of-market
Long-term hedge requirements	**Physical UCAP capacity** Forward purchase of physical UCAP	**Financial energy hedge** Forward financial purchase of energy (e.g., high-strike price call option)
Interconnection approach	**Current approach with CIRs** CIRs for firm resources in a structured "invest-and-connect" approach	**Flexible non-firm approach** ERCOT-style non-firm approach without CIRs, a faster "connect-and-manage" approach
Transition approach	**Phase in hedge requirements** Multi-year increase in capacity price cap; continue out-of-market procurements as needed	**Phase in hedges + enhanced ORDC** Multi-year decrease in capacity price cap; phase in energy scarcity pricing and hedges over time

Notes: *Strategic reserve was not explicitly included in PJM's Path C.

Our conclusions are that:

- **Either Path A or C may be able to break the credibility trap**, offering more efficiency and durability than the status quo, although no Path's success is clear or able to produce immediate results. Among the largest challenges, states would need to agree on hedging terms and scope, likely focused only on small customers in restructured states. Hedges would have to be phased in over multiple years to reduce the risk of broadly locking in high current prices that could later prove to be well above market.
- **Path A features lower implementation complexity** by retaining the current underlying market construct; but it raises concerns as a basis for very long-term hedges since the definition of "capacity" has and will likely continue to evolve over the life of a contract, for example by revising delivery periods, accreditation, and zonal representations.
- **Path C scores better for supporting long-term contracting and long-term efficiency**. As compared to capacity, energy forms a more concrete and enduring foundation for long-term contracting. Further, Path C offers economic efficiency improvements associated with "energy-only" markets: providing strong incentives for resources to be available when and where needed, attracting a portfolio of resources based on diverse market participants'

views of future resource capabilities, and avoiding distortions from the static and imprecise zonal treatment of locational capacity by instead using the nodal energy construct. The greatest benefits arise in combination with a more flexible, fast-track interconnection process for generation seeking non-firm energy-only interconnection service. This would allow new resources to enter more quickly and at lower cost and still access the full marginal value they contribute to avoiding shortages. Such a flexible system is well suited for a future with more fluctuating generation, storage, and usage patterns.

- Any **transition** from the current market design to Path A or Path C would have to be slow and deliberate. Both Paths involve large amounts of long-term hedges that must be layered in over several years to reduce the risk of locking in high current prices that could soon prove well above market, leading to dissatisfaction and possible risk of default. Path C would require additional time to develop the new energy market designs—and to let long-term hedges incorporate an understanding of the new design.

Table 1 below shows the summary of **our assessment of each Path** against the objectives of reliability, efficiency, implementation, construct durability, and alignment with state goals.

TABLE 1: SUMMARY ASSESSMENT OF PATHS

	Path A: Long-term Capacity Hedges	Path C: Long-term Energy Hedges
Viability & Durability	**Weaker fit for long-term contracting** • Implements long-term hedges, addressing credibility trap, but with uncertain feasibility and liquidity • Presents durability risks during the contract term as the capacity definition and underlying rules evolve (e.g., time periods, accreditation, deliverability)	**Strong foundation for long-term contracting** • Implements long-term hedges, addressing credibility trap, but with uncertain feasibility and liquidity • Uses energy as a more concrete, stable product for a long-term hedge
Reliability	**Maintains today's 1-in-10 approach** (but more supportive of investment) • Continues to procure target reserve margin via status quo capacity market, with price cap slowly rising • Supports performance via E&AS prices and CP penalties	**Achieves comparable reliability assurance, with upsides for adaptability and realized reliability** • Procures to the target via strategic reserve • Strengthens performance incentives through higher E&AS scarcity prices • Could relieve tightness through faster energy-only interconnection and more demand engagement
Economic Efficiency	**More administrative guide to investment** • Uses uniform capacity signals, accreditation, and deliverability rules • May better suit demand response under equilibrium conditions • Administrative rules can constrain ill-fitting resources • Existing CIR model slows efficient entry, unless reformed	**More efficient: rewards performance, lowers barriers, decentralizes investment decisions** • Strengthens incentives to perform during scarcity • Lets investors (and state planners) trade off firmness, speed, duration, and cost • Enables flexible demand to respond according to willingness to pay • Supports faster entry through non-firm interconnection
Alignment with State Goals	**Generally aligned with state goals for affordability and energy transition** • Can limit bill impacts by gradually increasing capacity prices • Can accommodate state procurements of preferred resources (albeit with reform potentially needed for greater decarbonization)	**Comparable alignment, with flexibility and interconnection advantages** • Can limit near-term bill impacts by gradually decreasing capacity prices while increasing energy scarcity pricing • Can accommodate state procurements • Requires less transmission for firm decarbonized resource mix
Practical Implement-ation	**Lower implementation burden, more readily fits current state frameworks** • Can leverage existing capacity market mechanisms for the new hedge procurement and allocation mechanism • Capacity is already embedded in retail rates, state procurement programs, etc.	**Substantial redesign required** • Relies on novel hedge procurement and allocation mechanisms • Also requires major changes to the underlying market construct: enhanced energy-market design, credit rules, strategic reserve, market power mitigation, and interconnection criteria

Finally, all market redesign efforts would benefit from many no-regrets enhancements to planning, operations, and markets—including several items identified by PJM—that could help mitigate supply scarcity. We do not address those in this paper but do address some in comments submitted separately regarding ways to facilitate the participation and contribution of all resource types.

II. Introduction and Context

1. Problem Statement

Customers across PJM face rising reliability risks due to the combination of rapid growth in large loads, retirement of existing generation, limited new entry, and significant challenges to the investment framework for attracting further investment.[2] Capacity prices increased sharply in 2025, causing a notable 2+ ¢/kWh increase in retail rates for customers that are not hedged against market prices.[3] Combined with other factors, the increase in retail rates driven by these capacity prices prompted a strong political response, leading to the negotiation and imposition of a temporary capacity price cap (called the "collar"). The collar limits capacity prices through May 2030 to levels below what would be necessary to support merchant investment in new generation. PJM articulates these challenges as the "credibility trap:" customers in restructured states rely on market prices for investment yet may again seek to intervene when those prices rise. The prospect of such intervention and the resulting impact on investors' perceptions of regulatory risk casts doubt on the durability of the market construct, undermining investment incentives and ultimately reliability.

Other stakeholders emphasize a different set of challenges, arguing that the rise in prices was not driven by market fundamentals, but by rapid data center load growth, combined with bottlenecks in permitting, interconnection queues, and electrical equipment. In such conditions, higher prices, even with longer-term contracts, will not yield new supply in the short term; said more simply, "ordinary consumers and their representatives do not want consumers

---

2 Reliability metrics (specifically for resource adequacy) are declining, with the capacity auctions for 2027/28 and 2028/29 clearing nearly 7,000 MW short of the requirement. Such shortfalls increase the risk of rolling outages, a risk shared by customers across the PJM footprint due to PJM's pooled structure.

3 For completely unhedged customers, capacity prices are passed through directly. The Base Residual Auction for the 2024/25 delivery year cleared with prices of approximately \$270/MW-day, about \$240/MW-day higher than that for the prior year. Assuming a profiled load factor of 40% for a typical residential customer that is fully exposed to each year's capacity price, their retail rate could increase by approximately \$240/MW-day * 100 ¢/\$ * 0.001kW/MW / 24 hr/day / 0.5 load factor = 2.5 ¢/kWh. Actual rate increases may differ.

to pay more for less reliability." [4] As a result, states proposed a lower price cap to avoid significantly raising rates during the period when a mismatch between new supply and large new demands could be expected. This formulation also rejects the idea that mandatory LSE hedging—the central element of PJM's solutions, as explored in the rest of this paper—would alone eliminate the cost impact. Instead, it focuses on Path B solutions to distribute and allocate the costs and reliability impacts of the ongoing resource adequacy shortfall, important topics which fall outside the scope of this paper.

While these important perspectives identify different framings of PJM's current challenges, long-term solutions ultimately must ensure the ongoing ability of restructured states to rely on merchant investment to replace retirements, meet organic (or all) load growth, and maintain resource adequacy. Achieving success in all of those dimensions will depend on PJM and states understanding each other's legitimate perspectives on the challenges jointly facing the region and working together to craft a durable solution.

Our aim with this paper is to provide decisionmakers with a deeper understanding of the possibilities of PJM's proposed Paths A and C to assist informed decision-making. We do so by elaborating on the critical elements that each Path would have to develop—a hedging construct for either Path, and for Path C several enabling reforms to the underlying spot market and resource adequacy framework—and then evaluating each Path against objectives. Our analysis draws on PJM's exploration of the issues surrounding each path, on others' relevant papers, and on observations from both PJM and other jurisdictions.

2. Paths Proposed by PJM

PJM asserts that long-term hedging could resolve the credibility trap and thereby stabilize the market-based investment engine that restructured states depend on for achieving resource adequacy. Mandatory long-term hedging requirements would protect retail customers (especially those in retail-choice states) from spot market price volatility, and thus attenuate the political response during high price periods of the business cycle. [5] However, such a structure would also support investment, both directly by offering greater revenue certainty

---

4 See Monitoring Analytics, Quarterly State of the Market Report Q1 (May 14, 2026), at 3. See also PA Office of Consumer Advocate, Revisiting Root Causes and Their Implications in PJM's Powering Reliability Through Market Design (July 15, 2026).

5 See, e.g., Borenstein, Customer Risk from Real-Time Retail Electricity Pricing: Bill Volatility and Hedgability (September 25, 2006). See also Wolak, Long-Term Resource Adequacy in Wholesale Electricity Markets with Significant Intermittent Renewables (February 18, 2022).

(mitigating both regulatory and market risk) for new physical generation that sells long-term hedges; and indirectly, by bolstering the credibility of the spot market as the basis for other, less capital-intensive, shorter-lived investment. PJM observes that long-term contracts have to align with spot markets (revised with higher price caps) in order not to distort market participants' incentives and to support liquidity for adjusting long-term positions.[6]

Path A: Hedging requirements under Path A would involve long-term contracts for capacity, in UCAP terms, compatible with the current market design. PJM sketches two potential approaches. One would require each LSE to enter the capacity auction with a large share of its forecast capacity obligation already hedged (illustratively 90%), supported by a forward resource plan and a substantial noncompliance penalty. The other would have PJM centrally procure a large share of capacity through laddered, multi-year contracts, leaving the capacity auction to procure the residual need. In either variant, PJM identified key outstanding design choices such as the required hedge level, forward period or contract term, and enforcement mechanism.

Path C: PJM's Path C envisions reliance on stronger signals in the energy market, both as a basis for long-term contracting for energy and in the spot market. This pathway represents a shift from the current market design, where the majority of scarcity value manifests in the capacity market, to a future where scarcity value is found primarily or entirely in the energy market, via operational scarcity pricing. This can be achieved by raising the energy market price cap and more accurately reflecting the value of operating reserve shortfalls in energy prices through the operating reserve demand curves, and (in the long-term) facilitating more demand-side participation to help set or affect prices in the energy market. To be distinct from Path A, we assume Path C would we replace the capacity market with a Swedish-type out-of-market strategic reserve mechanism to fill projected resource adequacy shortfalls relative to the 1-in-10 standard.

Path B: During supply shortages, PJM's Path B differentiates curtailment priority among PJM customers (either by region, customer class, or another metric), placing a higher degree of responsibility on loads to manage their resource adequacy. Because Path B alone is not comparable in scope to Paths A and C and addresses a separate problem statement and suite of issues, we do not address it in detail in this paper.

---

[6] On the other hand, if spot markets are not so aligned (with insufficient revenues to support investment), the primary purpose of long-term contracts can shift to directly stimulating new entry by providing adequate revenue.

3. Key Design Components

As shown in Figure 2 below, this paper specifies a core solution (and variants) for long-term contracting under Paths A and C, alongside the underlying market design for capacity and/or energy. We describe the key design components required to demonstrate the feasibility of Paths A and C, and we show the important contrasts between them. However, the paper leaves many other solution details for later design development stages. We focus on the following components in the paper and Figure 1; the body of the paper also explores key implementation details beyond the components specified in Figure 2.

- **Underlying Market Construct**: Path A builds from the status quo construct, with most value for addressing scarcity conditions reflected in the capacity market; Path C puts more or all of that value in the energy market, reducing or phasing out capacity market remuneration. Revisions to energy and reserve markets would allow energy prices to rise higher than today as shortage conditions tighten, but customer exposure to scarcity events would be limited by their hedges. Under Path C, residual risk might still be high on any customers with unusually high or unhedged loads, and on resources that have sold forward then fail. These risks likely call for a safeguard to reduce scarcity pricing under certain conditions.[7]
- **Form of Hedges Required**: Path A mandates long-term capacity hedges; Path C only long-term energy hedges. Authors such as Mays and Wolak suggest hedging of full-requirements services through Standardized Fixed-price Forward Contracts (SFPFC) to most fully and reliably transfer scarcity risks away from customers. However, SFPFCs include material amounts of long-term fixed-price energy transacted forward in non-scarcity periods, creating challenges for certain types of resources to supply such hedges including those with capabilities that poorly match the full load shape, such as storage. Or, gas-fired resources that can match the load shape but likely cannot hedge their fuel prices beyond about 5 years. We highlight the benefits and drawbacks of each approach.[8]
- **Scope of Hedging Mandate**: Any hedging mandate could be limited to residential and small commercial customers in retail choice states, to help protect the customers that are least likely to hedge themselves, and to avoid interfering with large customers to design their

---

[7] For example, ERCOT's "Low System-Wide Offer Cap" lowers maximum prices during years when the cumulative "Peaker Net Margin" has exceeded three times the cost of new entry. See [16 Tex. Admin. Code § 25.509 - Scarcity Pricing Mechanism for the Electric Reliability Council of Texas Power Region] at § (b)(6)(C) and (b)(6)(D).

[8] One way to set the strike price on a call option is to base it on the fuel cost of producing the contracted energy with an inefficient gas generator. The strike price will thus track the market price of gas. This is called a "high heat rate" call option since an inefficient gas generator has a definitionally high heat rate, meaning that it takes more gas to produce a single unit of energy than more efficient generators.

own hedging programs. For small customers to be sufficiently de-exposed from volatility to break the credibility trap, it is likely that a large fraction of their need would have to be hedged. The specific form and amount and timing of mandatory hedges would have to be established by the states and PJM. This is arguably the most challenging aspect of establishing the specifics for Paths A and C.

- **Hedge Procurement Process**: PJM would procure and allocate credits and costs to LSEs for residual needs of small customers, net of any state-mandated procurements. This structure would function similar to capacity costs today, which follow customers through LSEs based on their peak load contributions. If any state-mandated hedges are not of the same form as PJM's mandate, PJM would translate such hedges into terms that relate to the hedging requirements: under Path A, into UCAP terms via PJM-projected Effective Load-Carrying Capability (ELCC) ratings for each future delivery year (which retains the uncertainty and challenges associated with ELCC calculations); and under Path C, into hedge equivalents using analyses with some similarity to ELCC analysis. Such translations may be highly sensitive to PJM modeling assumptions.
- **Resource Adequacy (RA) Assurance**: Path A mechanisms would remain similar to today, with spot capacity auctions—whether prompt or 3-4 years forward—procuring the full RA requirement; any contracted pairings of resource+load passing through the auction as hedged. Path C in its simplest form is energy-only, as in ERCOT, where resource adequacy is determined by the market, and the traditional 1-in-10 loss-of-load expectation (LOLE) is not enforced or necessarily supported except by enriching the ORDC (Path "C3"). Given that this energy-only structure represents a significant departure from the current PJM market structure (introducing implementation complexity), we focus on Path "C1" with a traditional resource adequacy requirement supported by strategic reserve like that used in Sweden.[9] A strategic reserve would procure out-of-market additional resources as needed to meet projected shortfalls from the 1-in-10 standard. Resources in the strategic reserve would be available to run only in emergencies, and PJM would create energy price correction mechanisms to prevent their dispatch from depressing real-time energy prices and thus displacing in-market investment.[10] This approach avoids relying on a capacity market to meet such shortfalls, thus avoiding reintroducing capacity market complexities and uncertainties and the need to hedge that product (i.e., Path "C2," which reverts to Path A).

---

9 See European Union Agency for the Cooperation of Energy Regulators (ACER), Security of EU Electricity Supply 2025 Monitoring Report (November 20, 2025) at 11, 12, 24.

10 Similar to the Real-Time Reliability Deployment Price Adder (RTRDPA) ERCOT uses when deploying Emergency Response Service.

- **Transition onto the Path**: Each Path requires a careful assessment of transitioning to long-term hedges, particularly given the currently inflated cost of new entry. As a result, either Path A or Path C would carefully and slowly layer in hedges over time in order to avoid locking into large volumes of contracts at fixed prices that could diverge from spot markets as prices fall, leading to dissatisfaction and possible risk of default.[11] Path C has additional considerations transitioning away from the capacity market and into an enhanced energy market, as discussed below.

[11] The risk of default from customers would be low given PJM's authority to allocate FERC-approved costs to them; generators would need to demonstrate adequate credit to cover in the event that their physical plant fails.

FIGURE 2: CORE DESIGN AND VARIANTS FOR PATHS A AND C

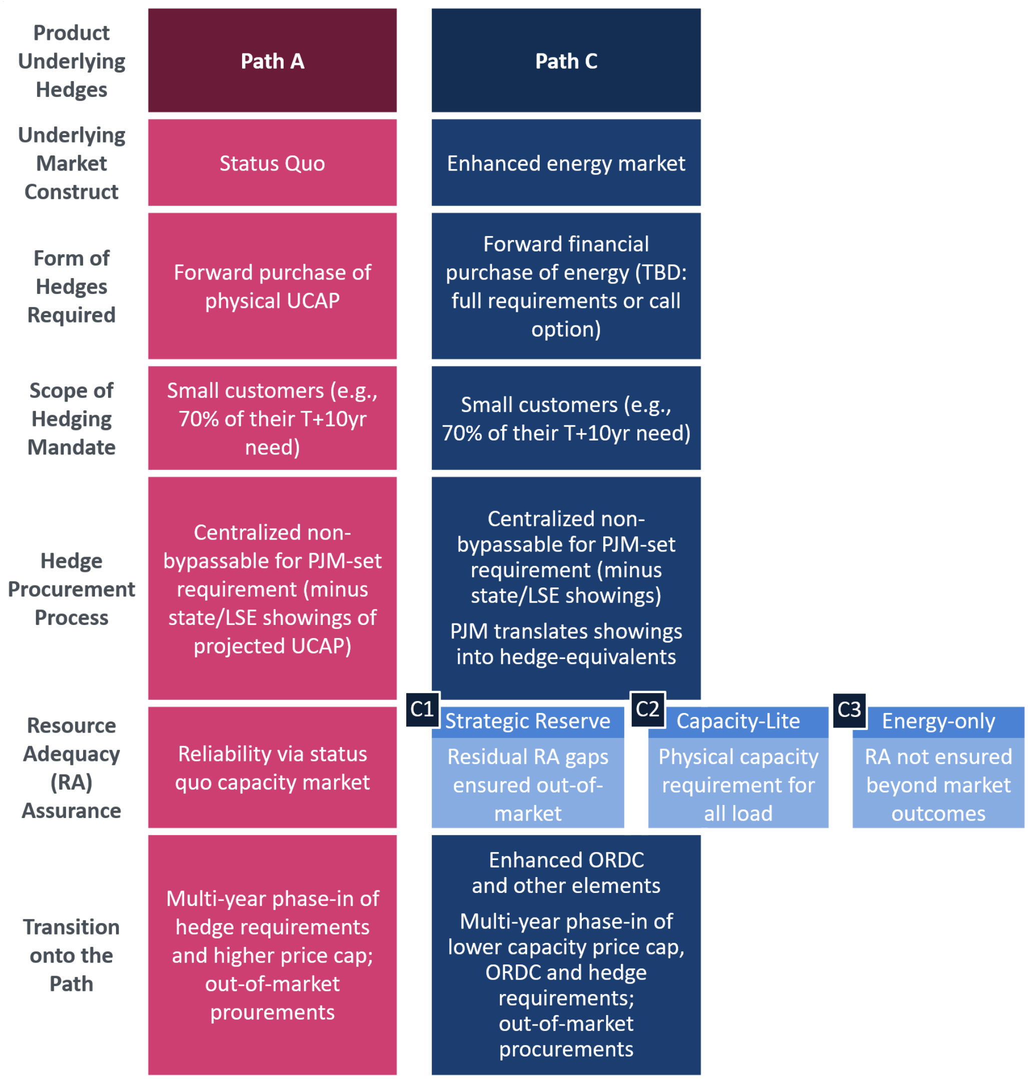

III. Development of Designs

This section first addresses the scope and size of any hedging mandate, arguably the largest challenges facing either Path A or C. Whether exposure to capacity in Path A or exposure to energy price spikes in Path C, much the same questions arise across the two Paths.

Then, separately for each Path, this section describes the underlying spot market construct, highlighting the elements that would differ from current arrangements and across Paths. We then discuss the associated hedging requirements structure, including the form of hedges, the hedge-procurement process, credit requirements for hedges, method for assuring RA, other implementation considerations, and the transition from current market arrangements.

A. Scope and Size of Hedging Mandate

We envision that either Path A or Path C would mandate hedging only for residential and small commercial customers in retail choice jurisdictions. Large C&I customers are likely to prefer to manage their own exposures themselves

For such customers, mandating hedges can de-expose them from price spikes, enabling them to pay a more stable price over time. The following example illustrates how hedges stabilize rates even when price spikes occur, although the example shows only a snapshot that does not demonstrate the reality that hedging cannot be expected to lower the average price paid over the longer term. Although this example is expressed in capacity terms, a similar example can be constructed for energy hedges under Path C.

- Suppose buyers have 70% coverage X years from the present to insulate against spikes in short-term capacity prices that are allowed to rise to an increased cap of $600/MW-day in present terms.[12]
- If the contract rate is say, $400/MW-day, and spot prices are also fluctuating closely around that level then suddenly spike to the cap, customer rates would increase as the unhedged 30% experiences the $200/MW-day spot-price increase, yielding a $60/MW-day increase overall.
- Under these circumstances, hedges held by customers limit their exposure to a quarter of the roughly +$240/MW-day increase from 2024/5 to 2025/6 that, combined with other factors, led to the current collaring of prices.

Among the largest outstanding questions on the scope of the mandate is how far forward to hedge, and what percentage of the forecast capacity requirement should be hedged. We do not craft a detailed proposal here, since those parameters would have to be established through extensive consultation with buyers and their representatives, and sellers. Buyers will have to specify and express their preferences; and any mandate should also consider the implications for sellers, given their own various resources' risks and hedging options.

---

[12] This illustrative cap is almost twice the current cap and roughly twice the estimates of "long-term Net CONE" included in The Brattle Group's 2025 CONE Study submitted as part of the Quadrennial Review. The estimate assumed that the cost of new generation reverts from current scarcity levels to levels before the recent increases driven by post-COVID supply chain constraints and especially demand-driven shortages of equipment, materials, and labor.

Regarding how far forward to hedge, buyers’ representatives would have to decide whether to hedge exposure to unexpected spikes that can arise in a single year (as from 2024/25 to 2025/26), which might call for 3-to-6-year terms; or whether to more ambitiously hedge against and smooth out the effects of multi-year business cycles (as in the current scarcity of combustion turbines and other equipment, labor, and installation contracting), which might call for at least 10-year terms.

Regarding the fraction of projected need to hedge and how that varies over time, the amount of hedges required should decline gradually into further-forward periods, so that customers are protected from near-term price shocks without being locked into large volumes of contracts that may become uneconomic as market conditions change. Further-forward periods are less necessary to hedge because long-forward prices are less volatile (so some quantity of forecast need can be procured the following year without prices rising as much as for delivery on shorter-forward timeframes). Full forward hedging is undesirable because the full projected need may not materialize if load growth lags expectations. And long-forward commitments can be costlier to sell from resources with limited remaining lifespans. Apart from those generic guidelines, establishing specific target percentages by year will be quite involved. Establishing any sort of demand elasticity would be even more involved.

Finally, it will be important to phase in hedges over time to avoid broadly locking in high current prices that could later prove to be well above market. A phase-in could ease into a laddering approach to purchase new contract tranches each year for future delivery, with each tranche adding to the coverage already secured for that delivery year. As a simple example, if the long-term hedging requirement is to be 70% of projected needs seven years forward, one could annually purchase 7-year delivery strips for 10% of the mandatory hedge amount. The full coverage would be layered in over the course of 7 years, followed by a steady state in which 10% is replaced in each year. This structure, similar to the overall structure of many existing default supplier auctions, would create a rolling pipeline of hedges rather than a single long-term procurement decision.

B. Path A

1. Underlying Market Construct

Under Path A, the Reliability Pricing Model (RPM) would continue to express demand—on a near-term basis, currently 3 years forward—for sufficient capacity to maintain current RA standards of one Loss of Load Event every 10 years (1-in-10 LOLE). We do not describe the fundamental framework of RPM here, since that is already well known. However, for spot auctions to be successful in procuring the full RA need in combination with proposed long-term forward procurements, several developments will be needed.

- First, incremental supply has to be available, and a separate paper we are developing explores enhancements to accelerate and expand supply participation with accurate accounting and compensation for their contributions.[13]
- Second, the spot price would have to be allowed to rise as needed to clear the market. Relatedly, to be efficient, market prices would have to be allowed to align in expectations (adjusted for risk) with forward prices, to avoid distorting retirement and entry decisions. That will involve a higher cap than the current collar, although proposed long-term hedges would smooth customers' exposure to price spikes).
- Third, revisions to energy and reserve pricing proposed by PJM through the Reserve Certainty Senior Task Force to improve operational reliability and efficiency could partially mitigate pressure on the capacity market. This is achieved by enhancing energy and reserve price formation, resulting in a larger share of generator revenues received through the energy market, even without transitioning to Path C.[14]

2. Design of Hedging Requirements

FORM OF HEDGE REQUIREMENTS

Full forward capacity contracts are agreements between a physically-backed seller of capacity and a buyer, with capacity transacted at a fixed price over some future period. These products would resemble the existing capacity commitments procured in today's 3-year forward RPM auction for 1-year delivery commitments, but with longer commitment terms. PJM describes illustratively 7-year terms for centrally-procured physical hedges under Path A;[15] we discuss term and quantities above in Section III.A.

One complication with longer-term requirements under Path A stems from a lack of clarity regarding resources' future ELCC at the time that a forward hedging requirement is imposed and contracts are executed ("ELCC risk"). However, to insulate buyers from short-term capacity

[13] See Pfeifenberger et al., The Brattle Group, Options to Address PJM's Near-term Resource Adequacy Challenges (forthcoming, 2026).

[14] See PJM, Powering Reliability Through Market Design: Addressing Rising Demand and Constrained Supply, and Stimulating Investment to Support Durable Reliability (May 6, 2026), at 47-48.

[15] *Id.* at 33.

prices, their hedging requirement would have to be specified in UCAP terms, consistent with the short-term requirement and markets. Such a structure would allocate ELCC risk to sellers, allowing fungible hedges and facilitating comparison of offers.

As discussed below, states that mandate contracting themselves, whether for the purposes of hedging or other policy objectives, could contract in any terms, such as tolling agreements in ICAP terms or clean energy/REC/capacity contracts in per-MWh terms, but PJM would still have to translate these into expected UCAP terms when determining how much residual need to buy itself.

HEDGE PROCUREMENT PROCESS

After the hedge requirement is established, PJM, if not some other centralized entity, could centrally procure hedges on behalf of load. States would be free to offset the requirement themselves, in whole or in part, through their own contracting or ownership. They would show their coverage to PJM, who would evaluate it in UCAP terms by applying its own estimates of the ELCC of underlying resources. PJM would then only procure any remaining necessary hedge volumes.

In retail-choice states, LSEs generally lack captive load, making long-term hedges difficult when costs decline because legacy above-market contracts can become stranded as customers switch to cheaper at-market suppliers. Centralized procurement by PJM can address this problem by allocating hedge costs and credits through non-bypassable charges, similar to today's capacity market. States could still procure long-term contracts for preferred resources through EDCs, with PJM procuring any residual UCAP requirement and allocating costs to LSEs based on customer peak load contributions. Such procurements could use auctions similar to current capacity auctions, though time-varying ELCC values would complicate auction design.[16]

The discussion in Section III.A above described fixed hedging requirements for each forward period, and laddered procurements of fixed amounts each trade year to meet to the total requirement. A variant is to have a sloped demand curve for procurements of multi-year hedges, similar to the VRR curve used in PJM's auctions for a single delivery year. While in theory this could incorporate graduated willingness-to-pay or beliefs about opportunities to

[16] For example, an auction could aim to buy, say 10% of projected UCAP needs for 7 years out, and that or more in the interim. Resources could offer on a $/ICAP basis, but be entered into the auction based on their time-varying ELCCs projected by PJM. One auction formulation could select the set of offers that minimize the NPV per UCAP MW-day (with future projected UCAPs discounted at the same rate as the dollars), although this is only one potential formulation.

procure hedges at lower cost on later trade dates, eliciting those from PJM and states may be less practical than agreeing on simple requirements. However, procuring to an inelastic (fixed) requirement has the challenge of introducing higher levels of price volatility and sensitivity than a sloped demand curve. This concern is partially mitigated in longer-dated forward markets for small slices of demand (e.g., 10% of need procured in a given auction) that are inherently more elastic and more structurally competitive than spot markets for full requirements. Perhaps the best approach would be to introduce demand elasticity using demand response concepts. States could identify a maximum price they would be willing to pay for hedges, such that for any price above that, they would commit to developing a corresponding amount of demand response. This could be implemented by their offering their own DR possibilities as supply (at a high price) in an auction to procure a fixed quantity of hedges.

CREDIT REQUIREMENTS FOR LONG-TERM CONTRACTS

Among the largest challenges with long-term commodity contracts is mitigating the risk of default for both buyers and sellers. Buyers can secure commitments through ties to customers, by having the contractual or regulatory right to allocate costs through non-bypassable charges. Indeed in this case, PJM would retain the right to pass through FERC-approved costs to LSEs; and to the extent that states are mandating their own hedges to avoid or reduce procurements on their behalf by PJM, they could show that they similarly had non-bypassable mechanisms to secure their own arrangements.

Sellers of hedges can back their commitments with physical capacity, although physical capacity can fail. For example, if a generator sells 10 years of capacity but suffers a catastrophic failure and retires after year 8, it would have to find and pay for replacement capacity to fulfill its commitments for years 9 and 10. Therefore, additional financial credit may be warranted even for physically-backed hedge sales. However, for each seller to fully financially back each physical MW sold would entail very large credit requirements, raising the costs of participating. There may be approaches where owners of larger fleets can self-insure; or PJM could qualify only capacity that can credibly show its prospects for operating over the entire period, then pool and socialize the risk of short suppliers defaulting. This is a topic for further development if PJM is to implement long-term hedging requirements, perhaps building on the credit requirement framework it already established for its backstop capacity auction underway this Fall.

RESOURCE ADEQUACY ASSURANCE

Path A mechanisms for resource adequacy would remain similar to today, with spot capacity auctions procuring the full RA requirement (and with any contracted pairing of resource+load associated with long-term contracts effectively passing through the auction hedged, as price-takers on both offsetting sides). Although a large fraction of the supplies would have been procured under long-term contracts, not all would. The remainder of the capacity needed to meet the RA requirement would be procured in the spot auction. Sellers might include any new or existing resources that have not already sold their capacity forward, although the shorter-term commitments may tend to favor less capital-intensive or shorter-term investments.

OTHER IMPLEMENTATION CONSIDERATIONS

We do not assess the full suite of implementation considerations for Path A. Instead, we note briefly only several aspects that must be carefully designed when implementing a feasible Path A solution: market power mitigation (evolved from today's physically-based mitigation framework), and cost allocation (at PJM – between LDAs; at the states – between retail customers).

TRANSITION ONTO THE PATH

Any transition onto Path A should layer in capacity hedges over a reasonable time horizon, accounting for today's shortage conditions. Similar to the transition to any future market design, PJM should evaluate ways of slowly reintroducing exposure to underlying price fundamentals to avoid introducing further challenges to market credibility and durability in the transition (or in an interim uncapped period). Once the contract volume is large enough to solve the credibility trap, the long-term spot market design corresponding to the "Paths" could be fully put in place, and the capacity price cap gradually increased.

C. Path B

Path B differentiates curtailment priority among PJM customers (either by region, customer class, or another metric) and may place responsibility on large new loads to manage their resource adequacy, perhaps with a requirement to bring their own *new* capacity. This approach does not address concerns about customers' exposure to spot prices through hedging, although it could reduce volatility for incumbent customers if it excluded from their market the effects of adding large and uncertain demand shocks to the system..

Path B contains numerous elements at the nexus of federal and state jurisdiction, including who has access to limited available resources, who manages hedging contracts and credit risk, and which customers ultimately pay the higher (hedged or spot) prices caused by the rapid growth of large loads. Ongoing conversations between the states and PJM are likely necessary to address these questions and ensure durable solutions to these challenges. Because the major elements of Path B are compatible with both Paths A and C, we do not directly address these important considerations in detail in this study. We note, however, that any changes in PJM's market design will also need to be coordinated with ongoing efforts by PJM and FERC related to the interconnection of flexible large loads and large loads that are co-located with new generation.[17]

D. Path C

While Path C relies on greater volatility in short-term energy spot prices, including higher energy price caps, customer exposure to this volatility can be hedged away. For example, residential retail customers in ERCOT (an energy-only market) typically choose fixed rates for energy going into each season or for up to two- or three-year terms, just as in PJM.[18] Indeed, the hedging requirements envisioned under Path C cover a defined high percentage of small consumers' load with long-term hedges, insulating these customers from volatility on the contracted volumes.

1. Underlying Market Construct

ENHANCED ORDC WITH HIGH ENERGY MARKET PRICE CAP

A central element of Path C is to express scarcity value and signal investment solely in the energy and ancillary services markets, rather than mainly in the capacity market. Path C thus requires energy market prices to reach either very high prices infrequently, or high prices somewhat more frequently. As the literature explains, the efficiency-maximizing way to effectuate energy market-based incentive signals (absent substantial volumes of bid-in demand elasticity) is with an ORDC with a high price cap expressing marginal system value.[19]

---

[17] See 195 FERC ¶ 61,209 (2026) and Docket No. EL-25-49-000 for FERC proceedings on co-location in PJM. See also 195 FERC ¶ 61,211 (2026) and Docket No. EL26-67-000 for FERC proceedings on accelerated large load integration in PJM.

[18] See, e.g., Public Utility Commission of Texas, Ways to Save: Shop for Electric Plans.

[19] See Chang, Aydin, The Brattle Group, Shortage Pricing in North American Wholesale Electricity Markets (January 26, 2018). Brattle studies have further shown that the exercise of market power is an unstable and undependable basis for expressing scarcity value in electricity markets. See Broehm, Chang, The Brattle Group, Market Power Screens and Mitigation Options for AESO Energy and Ancillary Services Markets (January 26, 2018).

Development of an enhanced ORDC could build on the solution package in PJM's Reserve Certainty Senior Task Force, which focuses on an ORDC structure to support short-term operational efficiency (assuming retention of the existing capacity product).[20] For the ORDC to further support RA, additional considerations include the (i) the scaling of marginal scarcity value curves, (ii) the maximum price, (iii) any RA-oriented uplift/extension, and (iv) circuit breakers:

- A high **scaling factor** for the ORDC based on the value of lost load (VOLL); this scales up the entire curve, following VOLL*POLL, the probability of lost load. ERCOT's IMM has recommended using a VOLL between $20,000/MWh and $30,000/MWh.[21]
- The **price cap** can be lower than VOLL to limit market participants' exposures on any short positions they may find themselves in, as hedges transfer exposure from customers to sellers. Higher values help attract investment and incentivize operational performance when most needed; they also increase efficient participation of price-responsive demand, as discussed below.
- Beyond the VOLL*POLL framework, including in the value of an incremental MWh of energy or reserves the avoided **cost of emergency actions** the system operator would otherwise have to employ.
- **Further upward or outward adjustments** (i.e., higher prices even for larger amounts of reserves) could be made to support a market equilibrium reserve margin closer to the traditional 1-in-10 standard, similar to the methods that have been used in ERCOT to adjust the curves for resource adequacy purposes.[22]
- **Circuit breakers** to limit market participants' exposures in extended extreme market conditions, without having to rely on prohibitively large credit requirements. This could be modeled on ERCOT's "Peaker Net Margin" (PNM) and "Low System-Wide Offer Cap" (LCAP), whereby if cumulative net revenues for a hypothetical combustion turbine dispatched virtually against real-time prices exceed 3xCONE within any given year, the energy price cap decreases from $5,000/MWh to $2,000/MWh.[23]

---

[20] See PJM, Reserve Certainty Senior Task Force Overview and Proposal First Read (August 19, 2026).

[21] See Potomac Economics, 2022 State of the Market Report for the ERCOT Electricity Markets (May, 2023) at 21.

[22] See Newell et al., The Brattle Group, Estimation of the Market Equilibrium and Economically Optimal Reserve Margins for the ERCOT Region (October 12, 2018).

[23] *Supra* note 7. See also ERCOT Nodal Protocols § 4.4.11.1. ERCOT maintains a similar circuit breaker, operating to reduce the energy market offer caps and resulting ORDC (to be based on a $2,000/MWh VOLL) once the PNM exceeds a defined threshold (currently $315,000/MW-yr). ERCOT publishes a running daily calculation of PNM for market participants, see ERCOT, Peaker Net Margin Webpage.

DEMAND-SIDE PARTICIPATION

While ORDC provides investment and operations signals based on administrative estimates of marginal system value, flexible loads that set or affect prices introduce information on the market value of power, reducing the importance of the administrative determinations. Infrequent but sufficiently high prices for investment signalling could occur under Path C largely via market bids. This would require the underlying market construct to allow flexible loads to set price via their bids to buy energy (while retaining existing supply-side energy market offer caps). It would also require a price cap above the short-run marginal cost of a significant portion of flexible demand—for example, PJM's illustrative value of $15,000/MWh, which likely exceeds willingness to pay of many flexible-capable loads.[24] Over time, those price signals could also cause the emergence of additional flexible loads and resulting reliability improvements.

To best support demand-side participation, Path C can build on PJM's existing Price Responsive Demand (PRD) construct to unlock additional flexible demand and enhance reliability. PJM's PRD mechanism enables specific loads to submit demand bids through a price-demand curve that commits the demand to reduce consumption as real-time prices rise.[25] PRD enables market clearing engines to consider and dispatch price-responsive consumption, introducing elasticity into otherwise largely firm (inelastic) demand. Properly designed, this approach can support substantial voluntary curtailment more flexibly than the current Load Management DR framework used for capacity market participation, which calls resources during emergencies without respect for individual users' time-varying value of energy (e.g., depending on where an industrial user is in its production process). Ultimately, enough PRD with a higher price cap could theoretically establish a new investment equilibrium between costly generation versus the frequency of curtailment of high-priced flexible demand (replacing today's capacity-market equilibrium between generation and load shed, and reducing the role of the ORDC in driving

---

[24] It is important to carefully consider the value used for the energy price cap. $15,000/MWh is relatively higher than PJM's current price cap of $3,700/MWh, and ERCOT's cap of $5,000/MWh. We note that recent surveys of Value of Lost Load indicate that many loads have an implied willingness to pay for electricity of less than $15,000/MWh, implying the ability to capture larger volumes of price-responsive flexibility. For example, for medium/large commercial and industrial customers in ERCOT, estimated average Value of Lost Load ranged from $6,507 to $12,783 for outages lasting more than one hour. See Gibbons, Sergici, The Brattle Group, Value of Lost Load Study for the ERCOT Region, (August 2024), p. 3. For AI data centers, "Compute Heat Rate" estimates for various workloads range from $1,465/MWh to $60,770/MWh, with a blended CHR (revenue-weighted average) of $8,000/MWh. See Royal, CHR Reference Values Q2 2026 (June, 2026). If data center willingness to pay considers marginal revenue without regard for sunk costs, the values would be substantially higher, but may still below $15,000/MWh for many workloads. It is possible that the value of compute workloads will differentiate further in the future.

[25] See PJM Manual 11 at §12.

investment signals). However, responsiveness in emergencies could also be made mandatory to satisfy a BYONC obligation, or under Path B-type state service differentiation framework.[26]

Not all demand participation would have to be dispatchable by PJM to contribute to reliability and market efficiency. For example, if a load with a reservation price of $1,000/MWh self-curtails when it sees prices rise above $1,000, doing so increases value for the customer, and it also helps to set prices lower and more reflective of marginal value, even if not with the same precision as a fully dispatchable load. This model may be more widespread than full dispatchability, especially in the near term. In ERCOT, few loads participate as "Controllable Load Resources" (CLR) that ERCOT dispatches as price-responsive demand and can set prices in real-time. ERCOT has recently observed nearly 3,200 MW of non-CLR flexible cryptomining loads responding to high prices and helping to ease shortages and contributing to price formation (compared with only a few hundred MW of registered CLR, also cryptomining facilities).[27]

The ORDC facilitates this kind of efficient participation in two ways: first, by making high prices predictably relate to the amount of operating reserves on this system, tying prices to market fundamentals that participants can aim to forecast; second, by configuring a gradual decline in prices as operating reserves increase, such that a single market participant's self-curtailment in response to prices should not entirely undo the high prices to which it is responding.

CREDIT AND RISK MANAGEMENT

Path C can concentrate large energy-market settlements (and settlements for ancillary services and FTRs) into short periods, increasing the need to mitigate risk that participants default on spot market transactions to cover short positions they have. This mitigation is achieved via increased credit and risk management procedures.[28] Credit and risk management procedures should not eliminate commercial risk but instead serve to protect the RTO and other market

---

26 We note that PJM has identified approximately 39 GW of ICAP resources covered by DOE's 202(c) authority, see PJM Presentation to Operating Committee, [January Cold Weather Operations]() (February 5, 2026) at 29. However, section 202(c) currently authorizes the Secretary of Energy to order only temporary actions; a Path C requirement directed at load would therefore need to rest on additional legal authority rather than on Section 202(c) alone, per 16 U.S.C. § 824a(c)(1).

27 In Summer 2025 for example, ERCOT observed that cryptomining loads reduced by a max of 3,151 MW during the system's peak day, plus an estimated 4CP reduction by other customers of 1,501 MW. See ERCOT, Item 12: Summer 2025 Operational and Market Review (September, 2025) at 10. In 2025, only about 240 MW of CLR participated in markets. See Potomac Economics. 2025 State of the Market Report for the ERCOT Electricity Markets (June 2026) at 14.

28 Credit and Risk Management considerations associated with the hedging requirements under Path C are addressed in the contract design implementation considerations subsection, below.

participants from defaults, and, ideally to encourage hedging (including through resource investment) by imposing accurate credit requirements commensurate with (particularly unhedged) settlement risk. Any Path C design would feature PJM in an expanded credit administration role, building on the detailed credit rules it administers in its current markets.[29] In addition, as noted above, a maximum-revenue circuit breaker triggering lower price caps would provide further protection from default by limiting total settlement exposure before a single event creates obligations too large for the market to fund.

Credit requirements apply to any entity that could find itself in a short position. Even generators can enter the spot market short if they have sold forward and their generator fails, creating exposure and leading to credit obligations. LSEs can become short to the extent that their load obligations exceed hedges (in ERCOT, for example, LSEs often hedge more than 100% of the forecast peak loads of their customers on fixed rates going into each summer and winter, yet ERCOT has to impose credit requirements to secure their ability to pay under an assumed set of adverse conditions).[30]

To ensure adequate credit and risk mitigation, PJM's risk assessment function might include the ability to assess the quality and terms of various risk mitigation measures. Said differently, PJM could conduct probabilistic market exposure assessments (not unlike current ELCC modeling) to determine the amount of risk "coverage" provided by each owned or contracted resource to accurately determine potential exposure and resulting credit requirements. These analyses would be based on PJM’s assessment of the probability of the seller meeting their exposure, considering elements like the generation resources backing the market position, the associated transmission contingencies and likelihood of curtailment, and the legal structure of any contract.[31] Alternatively, PJM could draw from ERCOT’s approach, which bases credit exposure on historical transaction volumes and forecasted future conditions,[32] or Australia’s approach in which the AEMO imposes credit requirements on each participant plus a “Prudential Margin,” which ensures no payment shortfalls in 98% of cases where a market participant defaults.[33]

---

[29] See PJM Settlement, Credit Overview and Supplement to the PJM Credit Risk Management Policy, v. 3.8 (July, 2025); PJM Tariff Attachment Q.

[30] See, e.g., ERCOT, Credit Management (September, 2025), at 54-60.

[31] *Infra* note 40.

[32] See, e.g., ERCOT, Credit Management (September, 2025).

[33] See AEMO, Report on the Effectiveness of Prudential Settings Methodology 2025 (December, 2025).

MARKET POWER MITIGATION

Path C's reliance on a higher energy price cap increases the importance of market power mitigation efforts in the short-term energy markets.[34] We envision retaining status quo supply side market power mitigation in the short-term markets. This includes the three-pivotal-supplier test for identifying structurally pivotal sellers, dispatching offer-capped resources using the lower of their cost-based or price-based schedules, requiring verification for market-based offers above $1,000/MWh, and capping the ability to set prices $2,000/MWh.[35] These mechanisms require careful review and consideration, particularly for instances where locations become structurally uncompetitive due to transmission constraints.

Creating the ability for demand-side resources to set price at levels above the current offer cap could introduce additional complexities. The possibility of price-responsive loads being part of an integrated portfolio (i.e., including generation and hedges) that can be net long at times raises questions about whether and how to mitigate demand-side bids. First, loads that simply respond to prices without submitting price-responsive bids, cannot conceivably exercise supplier market power nor be mitigated. Even for those that do submit price-responsive bids, mitigation may be impractical: loads' reservation prices are naturally variable and more irregular to characterize than a generator's economics grounded in a simple heat rate and fuel price. If mitigation is deemed necessary, mitigation mechanisms could be limited to portfolios with PRD that are net long on supply and that have structural market power (a potentially narrow category). Mitigation solutions could include limitations on bidding administration (e.g., requiring a load representative without any supply interests to make bids for loads) or on bid prices (e.g., restricting increases in bid prices relative to historical levels during periods of scarcity).

RELATIONSHIP TO GENERATOR INTERCONNECTION

Path C would eliminate the gateway of needing CIRs to provide value during expected scarcity conditions (value that currently resides in the capacity market). Instead, it would shift that value to the equal-access energy market, where all resources are dispatched economically regardless of CIRs versus energy-only interconnection service. As a result, new energy-only interconnections would have substantially the same value as CIR interconnections, and could access the market for scarcity value as much as incumbents. If combined with improvements to

[34] Market power mitigation considerations associated with the hedging requirements under Path C are addressed in the contract design implementation considerations subsection, below.

[35] See PJM Manual 11 § 2.3, 2.6.

PJM's interconnection process to make available and streamline a "connect-and-manage" style non-firm interconnection (as used in ERCOT), the shift away from CIRs presents a major opportunity to attract entry and interconnect it more quickly.[36] Path C plus interconnection reforms would enable resources to connect and access scarcity value (subject to congestion) without the delay associated with the construction of network upgrades to create their CIRs. Instead of PJM's minimum study requirements dictating the level of allowable curtailment (and requiring construction of network upgrades to mitigate the risk) a connect-and-manage framework under Path C replaces PJM's judgement with investor expectations of system availability.

Such an effort would require substantial changes at PJM. PJM's current energy-only interconnection service imposes relatively stringent requirements on developers compared to other regions.[37]Moreover, we expect any change to the CIR framework to be resisted by incumbents, particularly under Path C. Absent their ability to realize priority value in a capacity market, and finding themselves situated equally with new entrants with respect to a transformed RA value recognized only in the equal-access energy market, existing CIR holders would want to be compensated for the value they created through their funding of network upgrades. This might be addressed through allocation of Incremental Auction Revenue Rights (IARRs) or other ARRs. Or, they might want studies for energy-only interconnection service to be at least as stringent as current standards.

Apart from that special opportunity under Path C, generator interconnection processes could be improved under other paths too. Even today or under Path A, PJM could loosen its energy-only interconnection service criteria and provide a quicker "connect-and-manage" interconnection service as in ERCOT, and recognize such output in the capacity construct. Yet, compared to investors under Path C, PJM may be more reluctant to accredit generation under Path A when it is subject to deliverability risk. Indeed, PJM has resisted similar efforts to date, rejecting all output above firm deliverability levels in its reliability modeling (even for off-peak output that would be used to charge energy storage resources).[38] Path C offers the biggest

---

[36] The relevance of physical CIRs might be replaced by financial products, which could hedge sellers from delivery risk by providing the financial entitlement to the congestion over the pathway associated with their committed energy delivery. We do not assess the need for or design of these financial products in this study.

[37] See Pre-Workshop Comments and Exhibit of Tyler H. Norris of Duke University, Docket No. AD24-9, at 6-7.

[38] See PJM Interconnection, L.L.C., Capacity Interconnection Rights and the Accreditation of Generation Capacity Resources within PJM's Effective Load Carrying Capability ("ELCC") Construct, Docket No. ER23-1067-000 (Feb. 8, 2023), at 9, 10; "PJM proposes to modify its Accredited UCAP analysis in the ELCC model to cap the output of Variable Resources and Combination Resources in any hour at: (i) the resource's CIRs for hours in the months of

Continued on next page

impact of a connect-and-manage reform, since it would allow new entrants to realize more value from quick and non-firm interconnections.

2. Design of Hedging Requirements

FORM OF HEDGES

Several forms of potential Path C such hedges have been discussed by industry experts or noted in PJM's paper, including SFPFCs that lock in a load-shaped energy price in volumes that scale with actual load (i.e., fixed-price full requirements service) and financial call options at a high strike price or heat rate.[39]

SFPFCs would most completely de-expose customers from volatile shortage pricing, but they would also cover large amounts of energy throughout the year. This would raise the cost and the stakes of any hedging mandate for customers and for suppliers whose capabilities poorly match the full load shape, such as storage and intermittent resources. Furthermore, any requirement for hedges extending beyond about 5-6 years (as needed to de-expose customers enough from business-cycle volatility to break the "credibility trap") could become very expensive to provide. That is partly because gas-fired resources, which currently provide nearly half of the energy (and capacity) in PJM, can hedge their natural gas price exposure in forward markets with liquidity for only about 5-6 years forward.

A call option with a high strike price might better meet the goal of insulating customers from shortage prices without having to hedge against non-shortage prices and match the overall load shape. Call options are particularly well-suited to combustion turbines (CTs), whose economic value is concentrated in a small number of scarcity hours, consistent with the structure of options, which only require the resource to be available (and generating economically) during these hours. They also work for resources that can meet peak needs in addition to more routine energy needs, such as storage resources and baseload resources (nuclear, combined-cycle gas generators, etc.).

June through October and the following May of the Delivery Year, and (ii) the resource's winter deliverability MW . . . . By more tightly integrating CIRs and winter deliverability MW into the upstream Accredited UCAP process, PJM will derive a more accurate assessment of what ELCC Resources are capable of physically delivering . . . .". For order approving relevant tariff revisions, see 183 FERC ¶ 61,009 (2023).

39 See PJM, Powering Reliability Through Market Design: Addressing Rising Demand and Constrained Supply, and Stimulating Investment to Support Durable Reliability (May 6, 2026) at Appendix C.

Both instruments would be risky for several other resource types to provide, suggesting caution in requiring so many of them that the cost becomes very high. For example, wind and solar naturally would bear significant risk in selling options or SFPFCs, or any shape (i.e., of output over time) other than their actual output to-be-determined, the most common structure of today's renewable PPAs. Their ability to mitigate shape risk is improved if they are paired with storage or can be mitigated as part of a broader portfolio.

HEDGE PROCUREMENT PROCESS

As described above for Path A, states would be free to meet hedge mandates in whole or in part themselves through their own constructs, with showings of their coverage to PJM. Then PJM or some other centralized entity would procure the residual needs and allocate the hedges and their costs to LSEs during the delivery year.

States could also reduce the amount of hedges they have to buy, or limit the price they would have to pay, using demand response (similar to the corresponding discussion under Path A). States could commit to developing demand response as a physical substitute for paying more than a specified price for hedges.

One difference from Path A is that any showings of state-mandated contracts (procured outside the centralized procurement) would have to be translated into terms comparable to the specified hedge requirements, such that the residual requirement could be calculated. Suppose, for example, that that hedge requirement is in terms of energy call options with $1,000/MWh strike prices; and an LSE might have a tolling agreement with a gas-fired resource and a REC+energy contract with solar resources. For PJM to determine how many standard hedges it should buy, each contract would be evaluated by conducting probabilistic market value simulations of their output and loads alongside system conditions (similar to ELCC simulations, but with ORDCs and energy prices modeled instead of just MW of load shedding).[40] These models will indicate the amount of equivalent protection (e.g., how many call options) provided by the standardized hedge product, and therefore inform the amount of residual hedge procurement to be conducted by PJM.

RESOURCE ADEQUACY ASSURANCE

Path C could take one of three approaches to meeting typical 1-in-10 reliability standards.

[40] The PowerGEM SERVM model is set up to conduct Monte Carlo analyses with ORDCs and is used regularly by ERCOT, although not for the purpose described here.

- The first (which we call C1) would retain a reliability target—such as a one-in-ten loss-of-load expectation—and use a strategic reserve as a backstop. A strategic reserve is in place in Sweden, Finland, Germany, and Australia.[41] It consists of resources contracted outside the normal energy market and activated only when the market for installed resources enters shortfall over a period (such as a year). Likely candidates to provide strategic reserve include retiring generators and demand-response resources that would otherwise be uneconomic or absent from the energy market.[42] In order for the reserves to complement rather than displace in-market investment, energy prices would be adjusted if necessary to ensure that activation of strategic reserves does not depress prices. Possible downsides include (1) the possibility under structurally uncompetitive conditions that existing resources want to exit the market to join the reserve; (2) risk of slight inefficiency of reserving resources even whether their variable cost is lower than prices; and (3) investors' perceived regulatory risk that politicians will argue during prolonged shortages to release the reserves and not adjust prices back upward.
- A second option (C2) would combine a more energy-centered market with a smaller physical capacity obligation for all load, retaining a slimmed-down version of the current capacity market while moving most revenue recovery and contracting into the energy product. This structure eliminates the need for a separate strategic reserve, but retains many obligations associated with the capacity product. It essentially becomes equivalent to Path A with a higher ORDC, and less money in the capacity market.
- The third approach (C3) would eliminate a formal reliability standard and allow the market to determine the level of reliability, accepting potentially more load shedding while relying on mandatory curtailment by large loads to improve reliability as those loads become more common.

CREDIT REQUIREMENTS FOR LONG-TERM CONTRACTS

Credit must be managed for long-term financial hedge providers as well as for spot market participants under Path C. The considerations long-term hedges are largely the same as under Path A, discussed above.

---

[41] See Thompson, Spees, The Brattle Group, Clean Security of Supply in Europe (February, 2026) at 13.

[42] For context, AEMO procured 817 MW overall in two tranches of strategic reserve purchases in 2024-25, or approximately 2.4% of peak demand, the equivalent of roughly 4,000 MW when applied to PJM's 165,000 MW of peak load. See AEMO, FY2024-25 RERT Report (August, 2025).

TRANSITION ONTO THE PATH

As discussed under Path A, any long-term hedging requirements would have to be added slowly and progressively over time to avoid exposure to large volumes of hedges being out-of-the-money and inviting dissatisfaction and possible defaults.

Path C would involve the additional transitional step of designing an enhanced ORDC and associated market rules for circuit breakers, credit requirements, and market power mitigation as described above.

The transition from today's market to Path C could operate as a two-way ratchet: the capacity-market price collar would be maintained initially and lowered over time as energy-market price caps rise in parallel with the progressive layering of long-term hedges. The initial level of the energy-market price cap could be calibrated against the energy-market revenues needed to support investment while the capacity collar remains in place. The energy price cap could then rise only as hedge coverage increases and the market demonstrates that scarcity pricing can be absorbed without recreating the affordability, credit, and durability risks that Path C is intended to address.

IV. Assessment

We qualitatively assess and compare the Path A and Path C designs against the objectives of construct durability, alignment with state goals, practical implementation, reliability, and efficient resource entry and operations.

A. Viability and Durability

Path C has the advantage that its mandatory long-term hedges are based on a concrete product—energy—whose definition will not change over the many-year term of the contract (though its pricing may be affected by technology-neutral administrative changes to ORDC including the energy price cap). By contrast, Path A hedges against the more abstract capacity product, whose definition and value can shift more from administrative changes including: applicable time period (annual vs. seasonal vs. more granular), ELCC accreditation, CIRs and the study criteria used to assess deliverability, restrictive participation rules (e.g., for DR), the Variable Resource Requirement (VRR) curve, and changes to the ORDC that can inversely affect capacity value in the long run. Changes to the capacity definition could disrupt long-term contracts, either posing a risk to the durability of the market construct, or locking in capacity definitions that would struggle to adapt to potentially rapidly changing conditions (e.g., under decarbonization or a rapidly changing economy). An alternative variant under Path A would be to structure long-term contracts in more statically-defined terms, such as fixed $/kW ICAP payments for tolling rights, although that raises its own host of questions about how to define a hedging requirement and compliance, how to evaluate competing offers from different resource types (with varying multi-dimensional projected values), and how to incent availability at critical times absent seller exposure to capacity performance incentives.[43] Overall, Path A does not fit as well as Path C with a long-term hedging regime.

An additional set of issues, shared between Path A and Path C, arises in the availability, feasibility, and liquidity of long-term hedges.

B. Reliability

Path A maintains the current approach to ensuring reliability at 1-in-10 LOLE. The C1 variant of Path C uses an out-of-market strategic reserve to achieve this same outcome, and is the primary variant considered here. Path C2 achieves this using a more limited capacity market, retaining complexities associated with capacity while shifting focus to energy markets for entry signals, creating unique and potentially challenging overlapping complexities, which should be managed in the transition but would present challenges for a long-term market design. Path C3 removes the RA requirement and will result in whatever reserve margins the market delivers (likely fluctuating and lower even in equilibrium than the 1-in-10 standard, as ERCOT's "Market Equilibrium Reserve Margin" studies have explained).[44]

Path C might offer stronger operational incentives for resources to be available during critical periods than Path A, if Path C enables prices up to $10,000/MWh or $15,000/MWh in the most critical intervals while Path A retains incentives closer to today's $3,700/MWh energy price cap plus capacity performance incentives closer to $3,000/MWh. Path C's sharper operational signals can incentivize reliability-enhancing performance in all types of resources. However,

---

[43] For example, in Great Britain's capacity market, agreements can up to 15 delivery years and feature a fixed capacity obligation for each resource, meaning the resource's capacity value is also fixed over the full contract term. See UK Department for Energy Security and Net Zero, Informal Consolidation of Capacity Market Rules (July, 2026) at 43, 181-184.

[44] See, e.g., Newell, Carden et al., The Brattle Group, Estimation of the Market Equilibrium and Economically Optimal Reserve Margins for the ERCOT Region: 2018 Update (December 20, 2018); Newell, Wintermantel, The Brattle Group, Estimating the Economically Optimal Reserve Margin in ERCOT (January 31, 2014).

further inquiry would be needed into the implications of such concentrated signals on suppliers, as discussed during recent modifications to capacity performance incentives.

Regarding demand-side participation and performance, both Path A and Path C can incentivize flexible demand capability. The advantage of Path A hinges on the capacity market's ability to monetize participation by loads that prefer being paid for a capability in a predictable yearly revenue stream. However, under the current Capacity Performance construct (assumed to be continued under Path A), DR must provide unlimited availability, and the exposure to mandatory calls is unacceptable for some end-users, such as many data centers or manufacturers with batch processes, whose reservation prices can vary over time. Even for others, the cost of participating under Path A increases when the prospect of being called more often increases in a tighter market more reliant on demand response. In a future with more frequent calls, a voluntary energy-only approach under Path C may better allow participation by loads with high or variable reservation prices. Relatedly, customers with high-value reservation prices can avoid curtailment under Path C more often than lower-value customers, an efficiency that is effectively unavailable in Path A. Whichever path facilitates more demand-side participation has a reliability advantage.

Combining Path C with the more flexible and fast "connect-and-manage" approach to energy-only interconnection service described above could attract and interconnect new resources more quickly, offering both reliability value and economic value. In ERCOT, the energy-only and connect-and-manage approach (and other differences) connects new generation in approximately 20 months, about half as long as PJM's approximately 40 months, as shown in Figure 3 drawn from LBNL. Evidently, the ease and speed of interconnection and ability to equally access the market without firm injection rights more than outweighs the increased exposure they face from future entrants competing for the same scarce transmission until transmission capacity is enhanced. This allows the market to more quickly adapt to changing conditions and avoid potential reliability shortfalls attributed to slow availability of new supply.

FIGURE 3: TIME BETWEEN INTERCONNECTION REQUEST AND INTERCONNECTION AGREEMENT

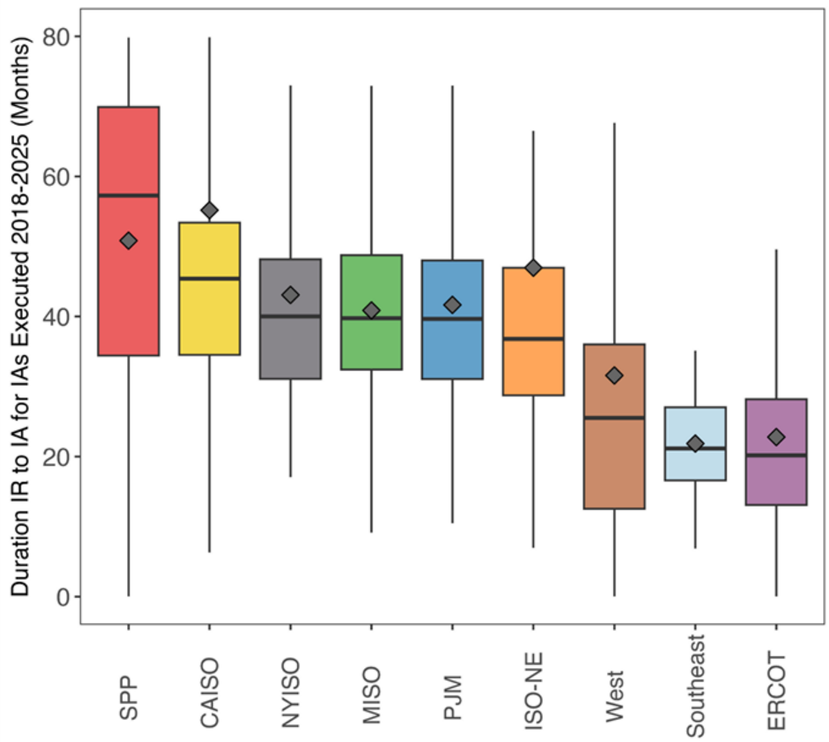

Source: Rand et al., Lawrence Berkely National Lab, [Queued Up: 2026 Edition](#) (June, 2026). Notes (from LBNL): (1) Sample includes 4,929 requests from 7 ISO/RTOs and 7 non-ISO balancing areas with executed interconnection agreements since 2004. (2) See prior slide for sample size by region. (3) A large portion of the 2024-2025 data in this analysis are from ERCOT. (4) Not all data used in this analysis are publicly available. (5) Date of IA execution for projects with IA agreement completed in 2023-2025 was not accessible in database format from several regions.

C. Economic Efficiency in Investments and Operation

Path C offers advantages that have long been articulated with markets that are more nearly "energy-only." First, by concentrating economic signals in actual operational shortages, it strongly incentivizes resources to be available when and where needed, by ultimately rewarding only actual performance. The market can respond accordingly, with any innovation, investment, or management practice that helps, including gas-fired plants winterizing and securing fuel, storage resources better optimizing their charge state, or all resources optimizing their maintenance schedules.

Second, it replaces a capacity-market system relying on highly structured interconnection criteria and capacity accreditations with financial signals that more fully privatize and distribute decisionmaking among market participants (or, to the extent applicable, state planners).[45] This can attract a portfolio of resources based on diverse views of which resource capabilities will best meet future scarcity patterns. It allows a diversity of levels of conservativism regarding asset investments and operations, and tradeoffs to be calculated by parties with a financial stake in the outcome, such as weighing faster/cheaper interconnection vs. firmer interconnection rights, or limited-duration capability vs. the potential for long-duration scarcity events, or consuming energy vs. curtailing it.

Third, Path C moves much of the necessary transmission investment from the interconnection process (firm network rights assignments to each resource) into regional transmission planning process to address emerging areas of congestion exhibited across all resources. Over time, this will reduce the amount of necessary network upgrades for interconnecting resources. This is enabled by allowing different resource types (wind, solar, storage, peakers, etc.) to use the same transmission capability at different times with reasonable amounts of congestion, rather than assigning a certain amount of firm transmission capability to each resource. Market participants will face more congestion risk, but this will be bounded by the regional transmission planning process (including market efficiency upgrades), and overall system efficiency will be higher as the transmission network can be smaller overall. An improved market efficiency construct coordinated with efficient long-term planning would replace many of these piecemeal network upgrades with coordinated and more cost-effective solutions responding in a holistic fashion to particular areas of system congestion together with other system needs. These benefits are particularly large for a decarbonizing grid with larger aggregate nameplate of resources combining solar, wind, storage, and firm generation.

Fourth, Path C can more efficiently orchestrate demand-side flexibility among consumers with diverse and time-varying willingness to pay for electricity. Demand-side capabilities are revealed by allowing energy prices to rise during scarcity conditions to levels that exceed the value of energy to many flexible consumers. Under those conditions, a flexible consumer will tend to curtail to avoid the high price (or to create value for contract providers). This contrasts with Path A, in which customers enroll in a PJM program that pays for curtailment mediated outside the energy market. Path C therefore offers several key advantages for demand-side flexibility: consumers can choose whether or not to curtail during any given high-price event; the attractiveness of flexibility does not decline under market conditions with more frequent or

[45] Even if PJM would still accredit diverse forward showings against a specified hedging requirement using its own probabilistic analyses, the spot market itself (with which forward markets would align) rewards only realized performance.

longer events; value is gained without the requirement to register or participate in a capacity market (avoiding transactional costs and potential penalties); flexible consumers (such as data centers) compete with each other to avoid curtailment, such that lower-value consumers are curtailed first; and (depending on how high prices may rise) it is possible to unlock higher volumes of demand side participation relative to Path A. PJM illustrates Path C with a price cap of $10,000 to 15,000/MWh, a value that is likely to exceed the willingness to pay of many flexible consumers.[46]

With such increased demand side participation, it is possible for reliability to efficiently increase above today's benchmark (a 1-in-10 LOLE) as entrants seek scarcity rents driven by consumer preferences rather than by administrative adders or mandatory procurements. With higher volumes of price-responsive demand setting higher prices during curtailment, the equilibrium resource investment can be set by tradeoffs between additional new generation (priced at the cost of new entry) versus additional voluntary load curtailment (priced up to $15,000/MWh), as opposed to trading off relative to load shed (conventionally priced at the Value of Lost Load). At that equilibrium, with sufficient price responsive demand at sufficiently high value, less new generation merely increases voluntary load curtailment, rather than increasing involuntary load shed. This construct also enables current Curtailment Service Providers (CSPs) to sell option contracts to translate their ability to curtail load into a stable revenue stream.

Path C can also accommodate more diverse business models by decentralizing decisionmaking and lowering barriers to entry. For instance, distributed resources under Path A must register with PJM in order to provide capacity value, with complex interactions and potential double counting if those resources are also behind the meter of a retail customer. This can pose a barrier to entry. By contrast, under Path C, infrequent scarcity value as well as daily energy value can be captured by reducing LSE consumption at LMP. This does not require registration or interconnection with PJM, thus avoiding potential barriers to entry. In a future with widescale proliferation of behind-the-meter batteries as well as controllable electric vehicle charging and electric heat, such facilitated market engagement can yield significant efficiency benefits.

D. Alignment with State Goals

States and their constituents may have many goals, but here we address energy policy goals regarding affordability and clean energy.

46 *Supra* note 24.

Regarding state's affordability goals, both Paths can also be transitioned gradually to limit near-term price impacts on small customers. Under Path A, the capacity price collar could be increased as long-term hedges are phased in, while under Path C the capacity collar could be reduced as energy-market reforms and corresponding energy hedges are introduced, helping shield residential customers from the higher and more volatile spot prices needed to support investment. In either case, the transition can therefore be paced to balance restoring investment incentives with states' objective of limiting customer bill impacts.

Regarding clean energy goals, both Paths can accommodate state procurement programs by crediting state procurements against the mandatory hedge requirement, preserving states' ability to pursue preferred resource types and contracting terms. However Path C has a number of features that could better accommodate a transition to more diverse and fluctuating generation, storage, and load patterns, through more flexible and efficient use of the grid.

In combination with an easier energy-only interconnection process, Path C can enable non-firm resources not only more quickly, but more efficiently than Path A. Path A may overbuild the grid by investing in transmission for every generator to be firmly deliverable under a fixed set of study conditions. Path C allows more varied use of the existing grid and rewards resources for whatever value they can provide in shortages that can arise under a range of conditions; then economic transmission planning would add transmission enhancements that are cost effective given an overall ensemble of grid usages.

Under Path C, a dynamic, nodal energy-only market replaces Path A's more static and imprecise zonal capacity market structure. Path A's zonal structure may become more imprecise as more variable and energy-limited resources are built and DR and DERs proliferate under an energy transition. Under these circumstances, system conditions fluctuate and differ further from the flow pattern assumed in defining single CETLs values for each zone in each year. With very different patterns of power transfers in different types of real-time operational conditions (e.g., low solar, low wind, peak load, long-duration, high thermal unavailability, etc.), inter-zonal transfer limits may need to be represented by a more granular model in space and time dimensions. Path C takes that accurately all the way to a nodal, interval-level model.

The current single-ELCC approach of the current capacity market design may not long work well with increasing penetration of resources with declining ELCCs. For example, the ELCC of 4-hour storage could be much lower in an LDA that is saturated with storage than it is at the RTO level, raising potential need for multiple nested ELCCs. This is evident in NYISO, which applies different ELCC ratings for the same class in different zones. The scope of this problem would grow as energy- and duration-limited resources increase.

E. Practical Implementation

At a high level, Path A solutions can be more easily implemented than alternatives by utilizing the existing capacity construct. Capacity is a core component of today's retail rates, default supply procurements in retail choice states, private contracts, market power mitigation, and cost allocation. Capacity cost allocation is designed to be portable to facilitate customer switching under retail choice, following customers based on their peak-load contributions. Existing mechanisms for 3-year forward procurement can also be readily adapted to longer-term contracts. With the centralized capacity market, the complex functions of resource adequacy planning are concentrated among the expert staff at PJM who prepare the inputs to the capacity market, providing analytical benefits for all PJM states and utilities. Capacity is therefore well established across the market and regulatory landscape and can be readily harnessed accordingly (notwithstanding durability shortcomings addressed above).

By contrast, implementation for Path C involves a significant overhaul of the existing underlying market design and associated processes for PJM, states, LSEs, customers, and others, including a new or adapted construct for allocating the credits and costs of mandatory hedges. Many wholesale contracts and retail rates can naturally transition to a world with lower (or zero) capacity prices and higher energy prices. However, no mechanism yet exists for procuring and allocating mandatory energy hedges, which (for the reasons we describe above) are likely to be structured as a central procurement (by PJM or others) with non-bypassable charges to applicable consumers via their LSEs. New market power mitigation measures may need to be developed, although the new design could likely leverage the existing ex ante and other mitigation in the energy market. And because Path C involves higher spot energy price volatility than Path A, new credit mechanisms would be needed to mitigate default risk. Further, Path C would require either abandoning existing RA requirements, or implementing supplemental capacity procurement mechanisms (via strategic reserve or otherwise) to continue to guarantee physical RA even in years with high investment costs.

V. Appendix: Responses to PJM's Questions

1. The path(s) (either as described in the PJM paper or alternative paths) that your organization believes PJM should pursue and why
   a. See §IV.
2. Additional or supplemental path(s) that you recommend States explore to support long-term market durability and/or customer cost stability
   a. See discussion of Path A and Path C transitions in §III.A.2 and §III.C.2.
3. Any specific areas of the market design that your organization believes should be included in the scope of the upcoming stakeholder process, including why focusing on those areas will help stabilize PJM's wholesale markets and support more durable resource adequacy in the region
   a. See discussion of Path A and Path C elements in §III.A and §III.C.
4. Areas that should remain out of scope in these stakeholder discussions and why it is appropriate to narrow the scope of work accordingly
   a. We do not address directly the scope or process of stakeholder discussions.
5. The process that your organization would recommend be used to structure the stakeholder engagement
   a. We do not address directly the scope or process of stakeholder discussions.
6. Any thoughts you have on issue prioritization and timing considerations for addressing these issues
   a. We do not address directly the scope or process of stakeholder discussions, but note important transition considerations in §III.A.2 and §III.C.2.
7. Any education or analysis that you would request from PJM to support the upcoming stakeholder engagement
   a. We do not address directly the scope or process of stakeholder discussions.