EKPC Backstop Concept Proposal

2/27/2026

Background

The PJM Board has directed staff to advance a backstop mechanism. Thirteen governors have committed via the January 16, 2025, "Statement of Principles Regarding PJM" to use state authority to allocate their fair share of cost of procuring needed new capacity to data centers and to protect residential customers. The concepts EKPC presents for consideration implement those policy aims by:

- requiring states to ensure Load Serving Entity (LSE) identification before any decision by PJM to procure additional capacity for them,
- sizing procurement narrowly around incremental need (future obligations that cannot be met by existing and planned RPM committed capacity resources) and
- by allocating costs to the LSEs whose incremental future growth-related short positions create the requirement.

EKPC emphasizes that it is not addressing the situation where an individual LSE may use market purchases instead of bilateral or owned capacity to meet *current* load needs; EKPC is focused solely on the i*ncremental future* load needs.

The purpose of this proposal is to outline a backstop structure that reasonably addresses need determination and cost allocation if PJM elects to centrally procure new planned capacity resources. The Reliability Backstop Action (RBA) is conceived as a one-time intervention designed to unstick the market and attract new entry that would not otherwise occur. It is not a substitute for the RPM. Nor should it disrupt, delay, or dilute incentives for projects already in development or for projects that would proceed absent a backstop. This point is fundamental.

The design and implementation of an RBA have significant implications for entities such as EKPC that are developing new generation to meet organic load growth. Load serving entities with demonstrated resource plans to meet customer demand should not bear costs for backstop resources that would duplicate existing or committed supplies. Moreover, planned projects undertaken by entities like EKPC are certain in a way that speculative market entry is not, and they must not be delayed or disrupted by an RBA process.

Summary of proposal

EKPC generally supports a capacity backstop construct in which PJM serves as a market facilitator, matching data-center buyers with new resource developers who then execute bilateral agreements. Such a backstop could adopt a standard contract structure, shared collateral arrangements, and an approach to address likely mismatches between the size of individual large loads and the optimal size of new generation projects.

Regardless of procurement mechanics, the design must define a procurement quantity and allocate associated costs in accordance with cost-causation. EKPC's proposal, therefore, focuses not on the backstop procurement mechanism, but on determining how much capacity is procured and how the procurement costs are allocated when procurement occurs on a pooled basis rather than through individual bilateral arrangements. EKPC's conceptual framework assumes PJM procures resources through a central auction mechanism, similar to the RPM-BRA.

The proposed framework requires PJM to verify incremental datacenter load growth by identifying the LSE responsible for serving each large load before that load is incorporated into the forecast. PJM would then establish the system backstop procurement target based on verified incremental datacenter growth and incremental organic native load growth that is not met by existing RPM-committed resources or by the LSEs' planned new resources, whether contracted bilaterally or self-built. Under this structure, PJM would set the procurement target at 98 percent of the verified aggregate pool net short, prioritize coverage of datacenter shortfalls in the clearing process, and fix that allocation at the time of procurement. Any remaining procured MW would be assigned to incremental native-load shortfalls, with costs allocated to the responsible LSEs. For clarity, EKPC proposes that the load obligations included in the backstop procurement target be strictly incremental to those already satisfied through the most recent BRA for the delivery year, as well as any future delivery-year projections PJM determines appropriate to include.

The proposal then assigns backstop procurement costs in direct proportion to the incremental needs that create the requirement: datacenter-driven quantities are charged to the LSEs serving those loads, and incremental native-load quantities are allocated to LSEs with corresponding shortfalls. In both full and partial procurement scenarios, datacenter obligations are satisfied first and fixed at clearing, while any remaining procured MW are applied to verified native-load growth. Settlement occurs at the LSE level in the delivery year, with cost exposure tied solely to incremental growth rather than total portfolio positions.

This structure is consistent with governors' commitments to protect residential customers and with the reasonable expectation that only large loads with verified suppliers / firm ESAs participate in the backstop.

Details of proposal

(1) Determining the regional capacity requirement

If the backstop procures the entire regional resource adequacy shortfall, rather than only the data center contribution, then determining the requirement becomes a system level incremental exercise. EKPC proposes that PJM quantify the procurement target from two components: the Large Load Adjustment associated with incremental data center growth and the incremental organic native load growth, each net of RPM committed existing or planned capacity resources including generation and demand response. At this stage PJM does not need to determine which LSE has rights to which resources. PJM only needs an accurate forecast of the incremental shortfall. In parallel, and to enable PJM to appropriately assign backstop procurement costs, each LSE must indicate, enforced by state action, if necessary, the large load(s) it will serve before that load is included in the forecast used to calculate the incremental shortfall.

This requirement ensures that there is a defined PJM customer responsible for any backstop cost allocation and prevents the default assumption that the host utility serves the load solely based on geography. It also reflects the fact that interconnection agreements are insufficient because they do not establish who bears the firm supply obligation for energy and capacity.

(2) Respecting existing procurement

A core principle is that the backstop should not interrupt or dilute credible in-flight activities that are bringing new resources to the market. EKPC therefore proposes that PJM recognize LSE actions such as constructing new generation or contracting with others to construct new generation and expand demand response programs. It also recognizes new generation projects may be directed by states in a manner that would provide certainty of their construction. This approach ensures that LSEs acting prudently to meet their expected future firm supply obligations to data centers or native load are fully credited and are not required to pay for committed backstop capacity.

(3) Sizing the backstop

PJM would then compute the aggregate net short by summing, across all LSEs (excluding FRR) in its footprint, the difference between obligations and committed resources for the delivery years in scope and set the backstop target at 98 percent (98%) of that amount. The two-percent margin recognizes forecast uncertainty, limits over-procurement, and preserves incentives for LSEs to deploy demand response or bilateral supply. If PJM centrally clears a backstop, it should consider a budget constraint on the procurement informed by VRR price caps for the quadrennial review to avoid extreme outcomes and deter opportunistic behavior under scarcity.

(4) Procurement priority and cost assignment

If the procured quantity meets the full backstop target, the cost assignment mirrors the incremental drivers of the need. The portion of the target attributable to datacenter load is allocated entirely to the LSEs serving those loads, while the portion attributable to incremental native-load growth is allocated to the LSEs with those shortfalls in proportion to their respective contributions. If procurement falls short of the target, the assignment first satisfies the datacenter-related quantity, which is fixed at clearing and is not reduced by any subsequent bilateral contracting by the LSEs serving those loads. Remaining megawatts are then applied to native-load shortfalls, but only to the extent those shortfalls reflect incremental organic growth included in the procurement target.

For example, if the target is 100 MW and 80 MW are attributable to datacenter growth, the full 80 MW is assigned to the LSEs serving those loads when 80 MW clear. If 90 MW clear, 80 MW are assigned to datacenter obligations and the remaining 10 MW are assigned to LSEs with incremental native-load shortfalls as measured in the Delivery Year. This structure is intended to encourage LSEs serving native load to take timely actions to address resource adequacy, as those actions improve system conditions in future years. In both full and partial procurement cases, the assignment follows the incremental need that created the requirement and does not involve any review of LSEs' full resource portfolios.

(5) Settlement and cost allocation

Settlement is at the LSE level in the relevant Delivery Year. Costs for capacity purchased to serve data-center obligations are charged directly to the LSEs serving that load, consistent with state commitments to protect residential customers and with the view that costs should track the end-use class that causes them. The balance, if any, is allocated across LSEs with incremental native load shortfalls in proportion to those shortfalls. Importantly, the allocation metric for native load is incremental exposure associated with the growth included in the procurement target; it is not based on each LSE's total capacity position. Only LSEs serving native/organic load can reduce their allocation exposure between clearing and delivery by taking actions such as demand response or additional bilateral procurement tied to that incremental growth, consistent with maintain incentives for LSEs serving native load to proactively manage their load obligation. Datacenter allocations, fixed at clearing, cannot be reduced via post-clearing bilateral actions.

Process flow

The proposed process is as follows: (i) PJM develops a system-level procurement target from LLA and native growth net of existing resources; (ii) states, to ensure appropriate cost allocation consistent with the January 16, 2026 Statement of Principles, ensure each large load has a designated LSE so that accountability and settlement can be enforced; (iii) PJM recognizes in-flight actions to bring new resources to meet load growth; (iv) PJM computes the aggregate net short and sets the target at 98% of that value; (v) PJM conducts procurement; (vi) PJM assigns procured MW first to data-center shortfalls (fixed at clearing), then to incremental native-load shortfalls; and (vii) PJM settles costs to the responsible LSEs accordingly.

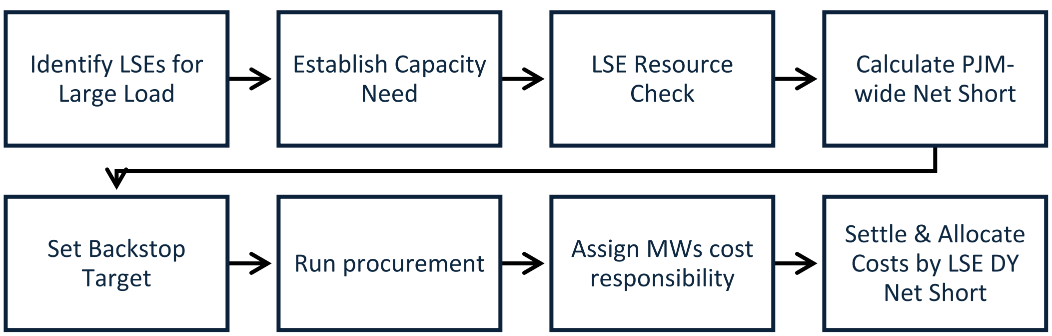

Addressing PJM's backstop design objectives

The framework addresses PJM's goals to protect residential customers, improve forecasting, provide predictable investment signals, and maintain a transitional posture that returns to market fundamentals once conditions stabilize. It is simple and narrow in scope; it transparently publishes the IRM gap, the aggregate net short, and the assignment by obligation type; and it clearly articulates a necessary state role. The approach also is consistent with large-load reasonable expectations that an executed power supply agreement (or ESA) is a commitment on the part of the LSE to provide the firm power supply.

PJM backstop design objectives	EKPC proposal meet objectives
Goals | *reliability, revenue certainty, protect residential customers, better forecasting, transitional posture*	LSE attribution associated with the incremental data center load growth forecast before considering it for inclusion in procurement target; priority assignment and direct cost charging to LSEs serving data centers; preservation of market incentives
Principles | *simple, narrow, transparent, stakeholder-driven, deliverability/risk aware*	Minimal scope limited to verified residual need; publication of IRM gap, net short, and MW assignment; explicit state role; deliverability risk reduced by limited volume and priority to datacenter shortfalls
Elements | *targets, eligibility, evaluation, term, obligations, deliverability, counterparty/allocation, transition*	Target = 98% of aggregate net short; LSE-specific assignment and settlement; one time procurement, settlement rules preserve (to the greatest extent possible) market exposure hedging incentives

Conclusion

As PJM considers options for the backstop, it should lean into cost causation, state authority, and protecting the integrity of active load serving obligation management driving new resources currently being undertaken by LSEs. The structure offered here protects native customers, fixes datacenter allocations at clearing, targets only the incremental need, and encourage incentives for LSEs to manage exposure through markets rather than administrative procurement. By tying cost allocation to the entities that cause the expected shortfalls incorporated into the backstop procurement target, the proposal preserves incentives for prudent forward procurement. Setting the target at ninety-eight percent of the point estimate of need recognizes the potential for forecasting bias and leaves room for LSE led demand response and other short-term market-based actions to address emerging shortfalls prior to the delivery year. The proposal assumes and requires that states will take the actions needed to ensure that the LSEs responsible for serving large loads identify themselves to PJM before that load is included in the backstop. This requirement is consistent with the expectations of large load customers that participation in PJM markets must be supported by a financial commitment and a power supply agreement (not just wires service agreement) with an accountable Load Serving Entity.

Appendix

Example of LSE incremental shortfall calculation

Assume LSE_A has forecast native-load growth included in PJM's forecast of *15 MW*. This is not the LSE's total load; it is only the incremental growth PJM included when constructing the Backstop Procurement Target. This approach is consistent with the design principle that shortfalls addressed through the backstop be tied to the causes of incremental system need.

Assume that that LSE_A has taken the following pre-Delivery Year actions to meet a portion of the forecast load growth: (1) a 4 MW PPA with a new generator for the DY; (2) 3 MW of additional demand response enrolled for the DY. These activities address 7 MW of the incremental load growth, thereby reducing the LSE's RBA exposure, consistent with the design principle that LSEs (only native-load LSEs) have the ability and incentive to address their incremental shortfall. This encourages such LSEs to actively manage their load obligation.

Incremental shortfall is calculated as:

Incremental Shortfall = max [Incremental Load Growth - Incremental Forward Actions - RPM-Committed Surplus ,0]

where:

- **Incremental load growth** is the load increase included in PJM's construction of the backstop procurement target
- **Incremental forward actions** are new PPAs, self-builds, demand response expansions, or other forward commitments recognized by PJM under the proposal
- **RPM-committed surplus** is the quantity of capacity cleared in the applicable (*line-in-the-sand*) RPM-BRA that **exceeds** the LSE's load obligation for the Delivery Year and is available to serve incremental load included in the target

Incremental Shortfall = 15 MW - 7 MW = 8 MW

In this example, the LSE did not have surplus RPM committed resources. Thus, the load quantity used to allocate any residual backstop capacity after the data-center allocations are fixed at clearing is *8 MW.*

Examples of backstop cost allocation

The following examples illustrate how the proposed backstop structure operates when:

(1) the full procurement target is procured
(2) procurement falls short of the target

In both cases, the incremental capacity need is mostly attributable to data-center load, consistent with current PJM large-load dynamics and the governors stated objectives to protect residential customers from data-center cost shifts.

Large-load verification is assumed to occur through state-enforced LSE identification, where necessary, , consistent with the framework described in the whitepaper and reasonable expectations that only verified loads with financial commitments and with power suppliers obligated to manage their energy and capacity needs should enter PJM's load forecast.

Example 1 | Full backstop target procured

Assume the RBA procurement target is 120 MW with the datacenter contribution equal to 90 MW, and the Native/organic load growth contribution equal to 30 MW. In this example, the RBA procures all 120 MW.

Assignment of Procured MW

- **Data-center allocation (fixed at clearing) |** Because the datacenter share of the incremental need is 90 MW, the entire 90 MW is assigned to the LSEs serving those data centers. This allocation is fixed once the backstop clears and is not reduced by bilateral procurement undertaken after clearing. This protects native customers from subsidizing large loads and prevents strategic behavior by datacenter LSEs after clearing.
- **Native-load allocation |** The remaining 30 MW corresponds to the incremental native load shortfall. It is allocated proportionally across LSEs based solely on their Incremental Shortfalls (where *Incremental Shortfall = Incremental Load Growth − Incremental Forward Actions*) not their full net-resource positions:

  - LSE_A incremental shortfall: 12 MW
  - LSE_B incremental shortfall: 10 MW
  - LSE_C incremental shortfall: 8 MW

  These LSEs receive precisely the MW reflecting their incremental contributions.

Cost Allocation

- Data-center LSEs pay for 90 MW
- Native-load LSEs pay for 12, 10, and 8 MW respectively

This example demonstrates the operating principle in a full-procurement scenario: each category of load receives exactly the MW it contributed to the incremental shortfall.

Example 2 — Partial Procurement: Backstop Does Not Meet Full Target

This example illustrates the shortfall scenario in which datacenter obligations are satisfied first up to the amount that clears, while native-load LSEs receive only the residual and bear costs in proportion to incremental native load shortfall.

As in example 1, assume the RBA procurement target is 120 MW with the datacenter contribution equal to 90 MW, and the Native/organic load growth contribution equal to 30 MW. In this example, the RBA procures only 100 MW, resulting in a 20 MW shortfall relative to target.

Assignment of Procured MW

- **Data-center allocation (priority and fixed) |** Because the datacenter share of the target is 90 MW, and procurement yields 100 MW, the first 90 MW of procurement is assigned to the LSEs serving data-center load. This fixed allocation ensures that datacenter-driven obligations are satisfied first and that native customers are shielded from costs caused by datacenter growth
- **Native-load allocation (residual and proportional) |** 10 MW remain for native-load allocation (100 MW procured − 90 MW data-center allocation). Native-load incremental need is 30 MW; therefore, only one-third of that need can be satisfied. The proportional allocation is as follows:

  - LSE-A: (12/30) × 10 = 4 MW
  - LSE-B: (10/30) × 10 = 3.33 MW
  - LSE-C: (8/30) × 10 = 2.67 MW

This allocation reflect each LSE's incremental exposure to native load growth. Recall that actions taken to mitigate this exposure (e.g., DR, new PPAs, imports) adjust these allocations in future Delivery Years to provide continued incentive to actively manage load obligations.

Cost Allocation

- Datacenter LSEs pay for 90 MW (full fixed share)
- Native-load LSEs pay for 4, 3.33, and 2.67 MW respectively