PROSPEROUS INFRASTRUCTURE

TECHNICAL VALUE STACKER NO. 1

THE JOURNEY TO BUILD FOR AI

PCLR Builds the Moat for AI Loads

TECHNICAL VALUE STACKER NO. 1 | LUMINARY STRATEGIES LLC | JULY 2026

LUMINARY STRATEGIES

ARUSHI SHARMA FRANK

ORCID: https://orcid.org/0009-0009-8087-1191

SSRN Author ID: 12311088

Prosperous infrastructure is infrastructure that compounds.

About the Series

Prosperous Infrastructure is a continuing series documenting the evolution of power grid market design in the age of artificial intelligence. The series chronicles first-of-their-kind technical, regulatory, commercial, and financial architectures conceived, championed, or propagated by Luminary Strategies LLC and/or the ventures it advises.

These papers examine how new infrastructure systems are designed, governed, financed, and operationalized—from market rules and transmission planning to private capital formation, software, hardware, and operational control systems. Collectively, they document the market architecture, regulatory frameworks, financing models, and operating systems that make prosperous infrastructure possible.

About Luminary Strategies

Luminary Strategies LLC (Luminary) is an infrastructure venture platform that builds technical policy while backing companies developing the next generation of power infrastructure. The firm works alongside and invests with entrepreneurs, and builds their partnerships with power utilities, grid operators, investors, and public institutions to co-create the technical, commercial, and regulatory foundations of critical infrastructure.

Luminary believes that prosperous infrastructure is infrastructure that compounds. Well-designed markets attract capital. Capital accelerates innovation. Innovation strengthens infrastructure. Stronger infrastructure creates new opportunities for investment, capability, and economic growth. The firm's work is dedicated to designing and accelerating these reinforcing cycles.

Luminary founder Arushi Sharma Frank is an Adjunct Lecturer at the Johns Hopkins University School of Advanced International Studies (SAIS), a Senior Associate (Non-resident) at the Center for Strategic and International Studies (CSIS), an Advisory Board Member of Conservative Coalition for Climate Solutions (C3 Solutions), and an Executive Council Member of the University of Texas Kay Bailey Hutchison Energy Center.

Abstract

Artificial-intelligence-driven load growth exposes a structural limitation in traditional transmission planning: large loads are treated as fully firm obligations even where portions of their consumption can be made visible, controllable, and responsive to real-time grid conditions. This is a static reality of power markets, agnostic to the characteristics of those loads and irrespective of whether industrial processes or artificial-intelligence layers sit behind the load fence line. This paper documents the development of ERCOT's Provisional Controllable Load Resource (PCLR) framework as the first large-load market architecture in U.S. energy markets to provide a technology-agnostic, freely scalable solution separating firm planning certainty from operational consumption for any load seeking to energize before planned transmission upgrades are complete.

The paper traces the design from two intertwined technical foundations. ERCOT-sponsored NPRR1188, filed in June 2023 and approved in November 2024, made individual Controllable Load Resources nodal market Resources whose consumption could be dispatched using locational shift factors, priced at a Resource Node, and represented through Energy Bid Curves. Luminary Strategies' 2025 transmission-planning discovery process exposed the missing bridge: a co-located battery could discharge while a large load continued consuming, yet the planning process lacked a replicable way to rely on that future operating behavior when deciding how much load could energize before permanent transmission upgrades were complete. Those two threads converged through Luminary Strategies' September 2025 Explanatory Comments white paper and Version Zero proposal filed with the ERCOT Large Load Working Group; the November 2025 filing and Planning Load Working Group presentation of Planning Guide Revision Request No. 134; its subsequent evolution through the Batch Zero stakeholder process; and final adoption within PGRR145 and NPRR1325 in June 2026. The resulting framework establishes Low Power Consumption as the load's grid-backed planning floor, i.e., firm power bounded by the amount ERCOT determines can be reliably served for the applicable year, and Maximum Power Consumption as the upper bound desired by the registering PCLR. The Maximum Power Consumption (MPC) provides the flexible (dispatchable), conditional operating range dispatched through ERCOT's Security-Constrained Economic Dispatch system. ERCOT sets the **maximum approved LPC** equal to the peak Demand amounts the study determines can be served, while the requesting load's **initial requested peak Demand becomes the approved MPC**, contingent on successful PCLR registration. The load can then accept those values or reduce either LPC or MPC in a new Form W, Part B; it cannot increase them above the ERCOT-communicated values. [26,30,3,4,5,1]

PCLR is a modern expression of Schweppian power economics. It reverses the anti-market drift toward administrative large-load curtailment, the latter which is characterized by opaque emergency calls, difficult-to-underwrite load shed, and mandatory interruption standing in for bid-based demand. PCLR puts the load where it belongs—inside nodal SCED, bidding to consume and dispatched as a Resource. The design enables earlier energization while preserving the transmission planning process for full firm service and shifting incremental delivery risk toward the private customer. It does so agnostically whether ERCOT receives ten such loads or one hundred. It sets a new bar for clearing supply and bid-in demand together through transmission-security-constrained economic dispatch, supported by the commitment processes that determine which capabilities are available to be dispatched.

The paper also examines the Adjusted Bid Cap, co-location and resource-attribution challenges, dispatch inside unit/resource commitment, the role of storage and onsite generation in private firming, the Value of

Compute Load as the economic basis for bid-in compute demand, and Grid 2.0 as a newly published open protocol standard for orchestrating assets beneath the grid-facing dispatch obligation.[21]

PCLR is a minimum sufficient product for turning large loads into utility-system-ready, dispatchable grid Resources. Its significance lies in establishing the institutional foundation for load-side connect and manage: the grid provides the firm floor, while the private sector underwrites and firms the rest. It is a prosperous infrastructure solution that compounds private returns and public goods: the grid grows alongside responsive private investment in safety, reliability, and planning certainty.

Prosperous infrastructure is infrastructure that compounds.

Executive Thesis

PCLR has two intertwined technical origins. The first was ERCOT's own move to take controllable load nodal. The second was a practical transmission-planning problem involving co-located storage and large load. PCLR became the market architecture that joined them.[26]

ERCOT filed NPRR1188 in June 2023 to correct the market treatment of individual Controllable Load Resources. Zonal dispatch factors could misrepresent how a large withdrawal affected a particular transmission constraint; ERCOT therefore proposed Resource Node treatment, locational nodal shift-factor dispatch, nodal pricing, Resource-specific Energy Bid Curves, and separately visible CLR consumption. By August 2024, some market participants were already asking whether that nodal treatment should extend to all large Loads, not only registered CLRs. The PUCT approved NPRR1188 on November 21, 2024.[26,30]

In 2025, multiple developers and ventures began working with Luminary Strategies on the second problem.[39] A battery discharging while a co-located large load continued consuming could reduce net grid withdrawal during the interval in which a thermal constraint mattered. The battery could be modeled and privately contracted for that beneficial operating condition, while planners still had to preserve conservative assumptions for firm service and for storage that was not itself a transmission asset. The planning rules did not yet provide a replicable, project-neutral path for relying on the future dispatchable behavior to determine how much load could energize before permanent upgrades were complete.

The two stories were the same problem viewed from opposite sides of the planning-operations seam. ERCOT had already created the future nodal operating Resource. Transmission planning still treated the same future customer as a static firm obligation. The insight that shaped the campaign was that planning certainty and operational consumption do not have to count on the same data sets. Planning can continue identifying the conservative amount of load supportable as firm service. Operations can manage a larger quantity of conditional consumption through nodal SCED. The customer receives speed to power by submitting to operational dispatch and privately underwriting the delivery risk above the firm floor.[26,3,4,23]

Luminary translated that insight through a sequence of technical work streams inside and outside ERCOT regulatory processes.[40] The August 2025 SCED paper supplied the operating thesis: dispatchable demand belongs inside the same security-constrained operating architecture that already observes the system, identifies constraints, clears Resources, and verifies performance: NPRR1188 had already supplied the market primitive and our proposal was contingent on its implementation. The September 2025 white paper then brought the two strands together, beginning from the proposition that NPRR1188 was essential to large-load operations and arguing that its incentive for nodal CLR participation had to come full circle into large-load interconnection studies. Version Zero of PGRR134 as discussed at the Large Load Working Group meeting converted that theory into proposed Planning Guide language; PGRR134 as-filed formalized the architecture; the November 18 PLWG presentations translated it into a developer-and-utility operating sequence. ERCOT then absorbed the core construct into the system-wide Batch Zero process through PGRR145 and NPRR1325.[2,26,3,4,5,24,25,1]

The resulting product is a legal and operational foothold from which the market can build solutions attributing the dispatchable characteristics of energy storage and energy shifting native to the load.[41] PCLR creates a market around private firming of power needs relative to grid constraints in operations, irrespective of how much physical transmission is ultimately planned and paid for. It does not prescribe one technology. It invites batteries, generation, workload controls, contracts, switching architectures, telemetry, and QSE orchestration to compete around an ERCOT-defined operating obligation.[23]

The architecture made simple:

ERCOT establishes the firm floor. SCED governs conditional consumption above it. The customer privately firms the band and bears the risk that remains to solve for peak constraints that would otherwise bind on the load and cause a curtailment event.[1,4,13]

Introduction

From Nodal CLR to PCLR

Artificial-intelligence-driven load growth made visible a structural limitation in traditional transmission planning in ERCOT: large loads were studied as firm obligations whose full requested withdrawal had to wait for the transmission system to be built around them. The operating grid, meanwhile, already had a five-minute engine for managing congestion, prices, ramp limits, outages, and Resource behavior. PCLR joined those two worlds without asking planning teams to relax risk-based studies of the steady-state power system, enabling earlier energization as a dynamic grid-utilization solution managed in operational dispatch.[1,2,6]

For purposes of this paper, speed to power means energizing before completion of the firm transmission upgrades identified for a load. A PCLR is the large-load interconnection application of ERCOT's existing Controllable Load Resource construct. ERCOT-sponsored NPRR1188, approved November 21, 2024, supplied the essential operating prerequisite: it moved individual CLRs from zonal dispatch and settlement to Resource Node treatment, nodal shift-factor dispatch and nodal pricing, replaced their existing Real-Time Market Energy Bids with Resource-specific Energy Bid Curves, and required the Resource's consumption to be separately visible so SCED could accurately reflect its effect on transmission constraints. PCLR then takes that nodal operating construct into the large-load interconnection process: the load registers as a Resource, accepts dispatch across an operating band, and consumes above its grid-backed planning floor only when ERCOT's network conditions permit. Planning certainty remains intact. Operational consumption becomes conditional, visible, and enforceable.[26,30,1,7]

Dispatchability—whether of generators, batteries, or consuming loads connected to the grid through site controls—is a feature of wires markets that can price physics and structure security-constrained commitment and dispatch around the physical limits of complex infrastructure. When the grid operator places the risk of injection and withdrawal on private resources instead of treating load as a static obligation, it privatizes delivery risk relative to available grid headroom and creates room for technologists and capital to do more. Grid planners continue doing exactly what they are supposed to do: study the load, identify the firm floor, resolve voltage and stability prerequisites, and plan the transmission required for full service under the economic and reliability tests governing prudent siting and spending. In ERCOT, SCED then manages thermal congestion as it appears in actual operations. That separation is the moat for powering artificial-intelligence inference and training faster on complex, multi-utility grids, and it is equally applicable to industrial load ramps, distributed UPS and backup systems, transitions from islanded microgrids to grid power, and fully dispatchable loads including bitcoin mining.[2,8,23]

The physical opportunity that made the planning gap concrete was simple: a battery or paired generator discharging while the load continued consuming. That is the operating condition in which private infrastructure can reduce net grid withdrawal during the same interval in which a thermal transmission constraint matters. The planning cases could model pieces of that behavior but did not provide a replicable pathway to rely on it for early load commissioning. By then, however, ERCOT had already approved the other half of the solution through NPRR1188: the future load itself could become a nodal, bid-in, dispatchable Resource. PGRR134 joined the two.[26,5]

PGRR134 made simple:
NPRR1188 made large controllable load nodal. PGRR134 made that nodal operating state usable for interconnection. ERCOT had already decided how the future Resource should operate; PGRR134 asked planning to recognize that enforceable operating state when deciding whether the load had to remain dark until every firm network upgrade was complete.[26,5]

Dispatchability, Not "Flexibility"

The distinction is easy to lose when "flexibility" becomes the organizing word. Flexibility can mean a tariff, a voluntary promise, an emergency curtailment right, a battery, a backup generator, a workload-management plan, or a modeled reduction that never becomes an operating obligation. PCLR is narrower and more specific. It is Resource registration, nodal telemetry, an Energy Bid Curve, ramp capability, SCED Base Points, and a floor the load must reach when the network requires it. Flexibility write-large, is about whether or not a load at its site control or loads behind its common point of coupling to the grid, are capable of controlled variation. That fact alone does not make the facility a dispatchable grid Resource. Flexibility describes motion; it does not describe the obligation or evidence that makes the motion useful to the grid. [2,26,10]

The distinction also explains why PCLR is a transmission-access product rather than a general large-load flexibility program. The problem being solved is not the abstract ability to reduce electricity use. It is the ability to energize before every firm network upgrade is complete while keeping the resulting transmission risk inside ERCOT's existing transmission security-constrained operating process. The customer's flexibility matters to the wires provider only to the extent that it produces the specific, measured nodal response the network needs. [3,4]

The Grid Provides the Floor

The grid provides shared infrastructure. The private project captures the economic upside of earlier access. PCLR ties those facts together. The public system continues to build transmission for firm service and broader regional value. The private actor assumes the interim congestion exposure and decides how much capital to place behind preserving consumption. This is the meaning of private underwriting in the PCLR context: dispatchability, control, capital, and firming are brought by the party seeking the incremental grid access. [13,20,23]

By implication then, the word "flexibility" collapses two different products. Baseline firming raises the level of load the grid is willing to treat as dependable, which is the provision of "flexibility" from paired and/or internal behavior of the load at site control boundaries connected to the grid which can be baked into the planner model and raise the firm power floor for the load. Operational headroom access is flexibility in motion with dynamic, wires-dependent reliability: this type of flexibility allows the customer to consume above that dependable floor when the system has room, subject to dispatch and a return-to-floor obligation. PCLR enables both: it creates the dispatchable band and motivates private investment in increasing the firm floor so that the dispatchable band becomes a bankable curtailment expectation. [23]

To be covered in more detail in future papers, a clear accounting layer still needed in the market, known as capacity attribution: proving that a private battery, generator, contract, or portfolio resource is actually firming the customer's band rather than merely helping the system generally. The financing problem lives in that band. The operator, customer, lender, and insurer need the same answer to four questions: what is firm, what is conditioned, what supports the conditioned megawatts, and what happens when that support fails. [23]

PCLR Origins in the Broader Chaos of Slow Grid Upgrades

Firm-load planning met its match when AI arrived: hundreds of commentators and analysts following the load interconnection freeze story began to pinpoint that it is visibly incomplete from the perspective of developers, but visibly precise from the perspective of system planners.

As the story goes, for decades large transmission-voltage customers were comparatively infrequent. A TSP could study one major load, identify the upgrades needed to serve it, commission the project in stages, and fold the energized facility into the background case before the next request materially changed the model. The process assumed that load was an obligation to serve and that the grid would eventually be built to support the requested withdrawal. [6]

Hyperscale development changed the problem's speed, size, and simultaneity. ERCOT became an early epicenter because Texas was already operating around a steady stream of refineries and large industrial-process customers, alongside the country's largest concentration of short-duration utility-scale battery storage. Beginning in 2023, hundreds of new load requests were changing the same planning cases at the same time. A load study could become stale before it was complete. Full requested demand across every project could imply transmission needs far beyond what could be permitted, financed, manufactured, and constructed on the commercial timeline of the load. The existing process had no durable capacity-reservation mechanism and no operating category between "full upgrades and permanent power are here" and "waiting for full power, and dispatchable in the interim". [6,17]

The commercial consequence was equally important. Large loads have intrinsic commercial value, and they need both speed to initial power and a credible schedule for reaching their expected power ramps over months or years. An AI campus, more so than other large loads, cannot rationally leave useful electrical headroom idle for five to seven years merely because the final network configuration required to serve its ultimate load has not yet arrived. The time value of access is immense. That value, however, does not entitle the customer to transfer incremental delivery risk to existing consumers on the power grid, particularly as all of those users are in line for firm power, and very few of them have the capital to privately back up their peak loads. The unresolved market-design problem was therefore layered: how can the grid permit earlier consumption, and who underwrites the interval between the power the grid can firmly promise and the power the customer wishes to consume while the transmission system catches up?[13,14]

Several early answers considered in the "flexibility marketplace" were attractive but incomplete. As early as February 2025, Luminary Strategies began its work with utilities and transmission planners to identify and cement a clear thesis: a generic election to be "flexible" does not define telemetry, ramp rate, duration, recovery, controllability, performance, or enforcement. Without these elements, "speed to power" simply means giving a load requesting faster grid access, more access than would otherwise be warranted to the parts of the transmission system that are otherwise preserved for handling emergencies that impact all grid users. It is an equally weak solution in a vertically integrated utility facing a binding transmission deliverability limit and in an organized market such as ERCOT, where the wires provider does not control regional operating headroom and cannot manufacture transmission capability through a bilateral flexibility contract. The clear answer everywhere is to have more transmission support assets funded by the load, or to have the load bring its own transmission support in the form of wires-friendly resource dispatchability. Depending on the part of the country, all that remains to be parsed is which is a better bit for the incumbent regulatory structure. For example, in ERCOT, transmission distribution utilities do not own the means of generation: here, storage providing transmission security is a market-based service delivered through constructs like ERCOT Energy Storage Resource, or, backing a PCLR. In Portland General Electric, where the utility may own or control the means of generation or transmission, a similar battery energy storage system can be owned, operated, or simply integrated by contract into utility operations like a virtual transmission upgrade, even if wholly owned by private third parties.

Static flexibility classifications make the same category error. They sort customers into administratively convenient buckets—fast, slow, short-duration, prolonged, critical, interruptible—without answering the operating question that determines whether another megawatt can actually flow: **what can this specific load do, at this node, against this constraint, in this interval, under dispatch?** Wires headroom is neither categorical nor static. It changes with topology, outages, generation dispatch, weather, contingency state, and the nodal sensitivity of each Resource to the constraint being managed. A preassigned flexibility tier therefore describes a characteristic of the customer. It does not establish a transmission-access right.

This distinction is familiar in transmission planning. A battery co-located with load may plainly reduce net withdrawal when it discharges, yet planners cannot simply convert the battery's nameplate capability into an equivalent reduction in firm load. The planning case must account for availability, state of charge, outages, common-mode failures, contingency performance, and whether the mitigating response is sufficiently enforceable for the system to rely upon it. Where a non-wires resource is directly incorporated into transmission planning, the reason is generally that its response has become a modeled and enforceable part of the reliability solution—with the telemetry, controls, operating authority, and performance obligations that status requires. A privately owned battery sitting behind or beside a customer does not acquire that treatment merely because its owner intends to use it helpfully: that is a truism in competitive energy markets as much as it is in vertically integrated areas permitting private third party deployment of batteries.

In fact, this is why behind-the-meter generation or batteries can serve a campus without itself creating a transmission-access right in ERCOT. Netting loads against co-located assets can improve settlement economics (treated in ERCOT systems like a Private Use Network) without changing the firm planning result or grid access timeline for the load. A bespoke Remedial Action Scheme (RAS) can solve a defined contingency without creating a freely scalable market product. Waiting for every permanent transmission upgrade preserves planning certainty by sacrificing the economic value of every unconstrained operating interval: this is the reality of grid planning-to-operations that governs load integration today.[10,17]

The architecture we wrote for speed-to-power for large loads therefore had to preserve the conservative planning answer, make the larger operating quantity conditional, and attach that condition to the actual dispatch engine. PCLR supporters did not ask for static planning models to become dynamic: we gave the dynamic question to the part of the power system already designed to answer it in SCED algorithms. Firm planning establishes what the grid can promise. SCED determines what the grid can support now.

Part I — 2023 to 2024: ERCOT Makes Controllable Load Nodal

Condition precedent for PCLR

Controllable Load Resources existed before PCLR. What changed in 2023 was ERCOT's conclusion that an individual CLR large enough to matter to transmission could no longer be represented through zonal approximations. ERCOT filed NPRR1188 on June 27, 2023 in response to the PUCT's Phase 1 market-design blueprint and its direction to increase the utilization of Load Resources for grid reliability and improve the transparency of demand-side price signals. [26]

ERCOT's stated reason was that zonal dispatch factors could misrepresent the effect of CLR withdrawal on a specific transmission constraint. Dispatching a Resource using its locational nodal shift factor was, in ERCOT's words, "essential for efficient congestion management." ERCOT cited roughly 17,000 MW of large-load interconnection requests then expected by 2026 and observed that some of those loads were interested in becoming CLRs; given their individual scale, nodal dispatch and settlement were "crucial for reliable operation of the grid." [26]

NPRR1188 did considerably more: the proposal assigned an individual CLR a Resource Node Settlement Point; used the shift factor of the CLR's actual nodal location for dispatch; settled the Resource's energy consumption at a nodal price; replaced the existing RTM Energy Bid with a Resource-specific Energy Bid Curve; enabled co-optimization of that curve with Ancillary Service offers in the Day-Ahead Market; restricted OUTL to a CLR that was truly outaged and consuming zero; and required the Resource's consumption to remain separately visible from other Load and generation at the site. [26]

The revision also arguably changed the conceptual definition of SCED itself. ERCOT's proposed language described desirable Generation Resource output using Energy Offer Curves and desirable CLR consumption using Energy Bid Curves, with both evaluated against State Estimator output, Resource limits, and transmission limits. The implication is clear: as made notorious through our firm's article, "Dispatchability Not Vibes," the demand side of the power grid is now supportable by technology which obviates the technical need to dispatch it solely as an emergency lever.[26] For the qualifying CLR, the quantity of consumption belonged inside the same security-constrained optimization. [26]

NPRR1188 comes with a small comment record: the central argument for nodal dispatch was largely uncontested. Lancium's October 2023 comments focused on the economics of co-location. It argued that a CLR should not be disadvantaged relative to other co-located Load by using total consumption rather than net grid Load for Load Ratio Share settlements, and that a rule which made CLR registration economically worse would obscure the true system cost of Load and create a perverse incentive not to become nodally dispatchable. [27] In April 2024, ERCOT agreed with the core co-location concern: co-located CLRs should receive the same treatment as other co-located Load. ERCOT separated that question from the broader policy debate over transmission cost allocation, keeping NPRR1188 focused on the market and operational efficiencies of nodal pricing, dispatch, and settlement. ERCOT also restored the ability of a CLR to update its Energy Bid Curve at any time, preserving the bid as an active economic representation rather than a frozen administrative curtailment value. [28]

Oncor's July 2024 comments were practical to the issues of wires owners: front-of-meter customer metering was straightforward and behind-the-meter CLR metering was not. Oncor identified TDU access, billing effects on other behind-the-meter entities, engineering limitations, and the need for site-specific agreement on the metering architecture, while preserving the utility's independent obligation to serve. The question raised in these comments was how to make nodal CLR participation workable, not whether nodal treatment itself had value. [29]

These comments defined a short procedural history. ERCOT's Protocol Revision Subcommittee initially tabled NPRR1188 for additional review and Large Flexible Load Task Force (later re-branded Large Load Working Group) discussion. By August 2024, ERCOT Protocol Revision Subcommittee (PRS) unanimously recommended approval, and the record notes that some participants supported expanding the proposed nodal treatment to all large Loads rather than only CLRs. ERCOT's higher review committee, the Technical Advisory Committee (TAC) then unanimously recommended approval; the ERCOT Board unanimously recommended approval; the Independent Market Monitor supported the rule; and the PUCT approved NPRR1188 on November 21, 2024. [30]

ERCOT had therefore solved the operational half of PCLR before PCLR existed. A large load capable of qualifying as a CLR could become a nodal Resource whose willingness to consume, physical location, and effect on transmission constraints were visible to the market. What ERCOT had not yet done was allow that operating capability to change when a new large load could energize, similar to how generators can consume on the system quickly through "connect and manage" without 100% peak demand-based deliverability to the grid. Luminary's intervention with PGRR134 quickly catalyzes the interconnection bridge. [26,30]

Part II — May and June 2025: A Missing Case Study

ERCOT's existing BESS and load cases

In May 2025, Luminary's work on co-located large loads and Battery Energy Storage Systems converted from a project study question into a market-design inquiry. In correspondence discussions, ERCOT staff explained that a co-located Energy Storage Resource (ESR) intending to use its full charging and discharging capability would be studied through three fixed cases:

- maximum ESR discharge while the load is not consuming;
- maximum ESR charging while the load consumes at maximum demand; and
- the ESR off while the load consumes at maximum demand.

The missing case was the one that mattered for speed to power: ESR discharge while load consumption continued. That case represented the condition in which a private asset could reduce the incremental transmission burden created by the new load without shutting down the customer's operations.

Relatedly, the distinction between thermal and non-thermal limitations—and how each is handled through planned upgrades and interconnection prerequisites—disciplined the solution from the beginning. ERCOT SCED can economically redispatch Resources to manage modeled transmission constraints whose relief is represented in the real-time network model. It is not a substitute for the interconnecting utility's voltage studies, physical interconnection facilities, protection design, short-circuit requirements, or stability prerequisites. A project still needs a safe electrical path to withdraw power. Under ERCOT Planning Guide §§ 5.3.5, 9.2.5, 9.5.2, and 9.6, the applicable stability, system-protection, interconnection-equipment, telemetry, modeling, and energization conditions must be satisfied before Initial Energization, irrespective of whether thermal headroom may later vary in operations to the advantage of an energized load. [8,31]

This boundary mattered politically as well as technically. PGRR134 did not ask ERCOT to certify a battery as permanent transmission infrastructure, compel a wires company to create a non-standard, private control strategy, or reopen planning methodology on the Batch Zero timetable. It asked the planning process to preserve the firm result and the operating process to recognize a future Resource that would be dispatched under existing market discipline. The market could move on this ask because it was shaped to move: PGRR134 shaped innovation into a complex planning architecture which carefully managed a dozen moving parts.

Netting for operational visibility versus transmission planning

In 2025, Luminary asked whether partial Wholesale Storage Load netting or a future Netted Network construct could change planning treatment. The question was narrow: if settlement rules recognized the battery's charging and discharge behavior differently, could the utility study ESR charging load as non-firm and credit dispatchability in transmission planning? ERCOT staff's answer was that planning did not then contain a non-firm treatment of load, so a metering or settlement change would not deliver the intended planning benefit by itself.

That answer moved the design work away from netting as the central speed-to-power solution. (It remains however, a critical grid operations visibility solution and will be covered by subsequent papers). Private Use Networks and other netted arrangements can create real commercial value for co-located generation and storage, for example, enabling co-located batteries to receive Wholesale Storage Load treatment in settlements. They do not create a general rule allowing the load to energize before the firm upgrades required by the load study. Netting solves visibility and market accounting problems, and we were attempting to solve transmission access problems.

This moment was an important intellectual fork. A netted site can look smaller to the grid when onsite supply is operating. The same site can look abruptly larger when that supply trips, fuel becomes unavailable, a common transformer fails, or a battery reaches its duration limit. All of these limitations from the concept of netting without controllable loads, have since become clearer in ERCOT's attempts to implement its new Withdrawal Limited Private Use Network construct. Planning still needs the gross load, the resource configuration, the largest internal contingency, and the automatic response that keeps withdrawal inside the authorized boundary. Netting changes the accounting lens. It does not eliminate the need for a controllable load or an enforceable fallback. [17]

PCLR therefore grew from the load side outward, and our efforts clarified quickly in 2025 to focus on it and away from netting. The load would become the resource ERCOT could always call, and private parties could develop the insurance policy needed to make that call bankable at the site control. Batteries and generators would remain separately visible where required and could protect the customer economically and physically, if co-location was intended. The site's private control architecture would decide which asset moved first. ERCOT would retain a direct string to the consuming Resource.

As Luminary canvassed ERCOT staff and multiple transmission planning organizations, whether or not there was study separation of loads and batteries, or loads and generators submitted together in an interconnection request, turned out to be less consequential than initially appeared. Load and generation interconnection studies answered different questions and remained in different planning workflows, but each study still accounted for the presence of the other facility. A utility could model the battery reducing thermal loading. The unresolved question was whether ERCOT's Planning Guide and Protocols allowed that modeled behavior to govern the amount of load commissioned before permanent transmission upgrades were complete.

The engineering capability also existed, though standardized guidance from ERCOT was reported as lacking by multiple grid planners at the wires companies. These companies reported that they needed ERCOT guidance before treating battery or generator response as enforceable mitigation for a load-study violation. NPRR1188 changed what the answer could be: the planner did not need to bind the private battery itself into the load study. ERCOT could instead bind the future load as a nodal CLR through SCED and let the customer choose how privately to preserve consumption when dispatch arrived. [26]

Battery first; load second ("just in time, last in line")

The technical discussions to this point had produced the first private-firming sequence that was worth enabling in permanent market design. With the ideal solution, energy storage would respond first because storage preserves customer operations. Onsite generation could extend that response when duration ran out. Native load dispatch remained available as the final layer. A developer sees the battery or generator as the least disruptive response, as well as the most capable of being hedged relative to nodal congestion exposure and system-wide risk. The grid sees the load itself as the cleanest decrement to withdrawal, essential to grid planners as loads get larger and more capable of disrupting larger power system operations. PCLR eventually gave both views a place in the same architecture.

Dispatchable load became the obvious operational answer to the planning constraint. Where the planner could not bind a battery commitment into a static load-study case, ERCOT could bind the load itself through SCED. NPRR1188 had already supplied the Resource Node, Energy Bid Curve, nodal shift factor, and dispatch obligation. The private site could then decide how much battery, generation, switching, and workload control it needed to preserve its own operations while the registered load delivered the required network response. [2,26]

A note on PUNs, RAS, and SPS

ERCOT already allowed generators and loads to operate together in Private Use Networks. Those arrangements remain important for co-location, settlement, self-supply, and commercial contracting. Their value did not support the speed-to-power problem because the battery or generator was not being credited as dynamic mitigation that changed the load's firm commissioning result.

Remedial Action Schemes and Special Protection Schemes can solve specific contingencies through engineered triggers and defined corrective actions. Their bespoke nature makes them a poor general market product for hundreds of potential large loads. PCLR took a different path: the load becomes the dispatchable resource, while project-specific electrical design remains private and technology-neutral. Minimum interconnection facilities, special protection, voltage resolution, and stability prerequisites stay with the TSP and ERCOT. Consumption above the firm floor enters SCED only after those gates are satisfied.

Design insight shapes as "bring your own everything"

By the end of the first phase of work, the problem statement had changed. The question was no longer whether a battery could be modeled next to a load. The question was whether planning could establish one quantity for certainty while operations governed another quantity for consumption. Once framed that way, the solution became scalable and technology-neutral. Any combination of storage, generation, software, workload control, contracting, or native load reduction could support the customer's preferred operating outcome. The grid needed one enforceable result: the registered load follows SCED.[3]

The breakthrough baked into PGRR134

Planning certainty ≠ operational consumption.

The former establishes the firm floor. The latter can move within a dispatchable band. [3,4]

Part III — August to November 2025: Two Architectures Converge

Grid-visible and dispatchable

On August 2, 2025, Luminary published a broader operating thesis: large flexible loads able to declare, measure, and deliver dispatchable adjustments are not exotic edge cases. State estimation establishes the system picture. Contingency analysis identifies the risk. SCED or its functional equivalent determines the secure economic dispatch. Telemetry and performance evidence establish whether the response can be trusted. Flexibility belongs in the normal dispatch continuum before the grid reaches firm load shed. [2,12]

That work supplied the effort leading up to PCLR. NPRR1188 had already done some of the lifting, though none of it was yet implemented. The August paper explained why a nodal, telemetered large load belonged inside the ordinary security-constrained operating continuum rather than outside it as an emergency exception. ERCOT did not need a new physics engine. The remaining design question was how to make that existing operating capability usable by transmission planning and large-load interconnection. [2,26]

The two architectures converge

By August 2025, ERCOT had already determined through NPRR1188 that an individual CLR's location and megawatt movement had to be represented nodally because those movements could materially help one transmission constraint and harm another. Luminary's storage-and-load work had separately established that the planning process lacked a replicable way to rely on future private operating behavior to accelerate energization. [26]

The missing step was to make that future nodal operating state legible to the interconnection study. The grid did not need to forecast every future Base Point or prescribe the customer's battery, generator, or workload response: that task is as unnecessary as it is impossible to deliver at scale and through open, competitive means. The grid needs an enforceable floor to which the future Resource could be dispatched when a studied thermal violation bound. [26,3,4]

September 2025 was the point at which the two work streams became one converging architecture. Luminary's white paper presented at LLWG (the "Explanatory Comments") began from the proposition that "NPRR1188 is an Essential Solution for ERCOT Operation of Large Loads" and argued that the incentive for loads to participate as nodal CLRs had to come full circle into large-load interconnection studies. If ERCOT would operate the future customer as a nodal Resource capable of being dispatched to solve constraints, the planning process should recognize that enforceable future state rather than model the customer solely as an unconditional firm withdrawal. [3,20,26]

A virtuous circle for load flexibility and reliability planning

NPRR1188 makes flexible load legible to operations. PGRR134 makes flexible operations legible to firm planning.[26,5]

Texas bitcoin miners supplied important operating evidence which Luminary relied on to discuss how this proposal would be different than prior constructs which enabled loads to simply "opt out" of dispatch. Bitcoin as an industry showed that large computing loads could register as Controllable Load Resources, respond rapidly to prices and reliability conditions and follow primary frequency response, and reduce consumption without the process damage or restart penalties associated with many industrial loads. ERCOT already knew that compute could move, and the goal of NPRR1188 was to make that movement bankable for transmission security. [22] PCLR itself is a massive opportunity for cryptocurrency mining and other highly flexible assets to become the most dispatchable class of loads to support grid reliability: PGRR134 converted a known operating capability into interconnection currency which can be utilized by any large flexible loads up to the task of responding to binding grid instructions and saving the system well before emergency conditions force mandatory shutoffs for Texas power consumers. [3,4,5,22]

Version Zero and the explanatory comments/white paper

On September 18, 2025, Luminary and Agentic Infrastructure placed the architecture into the ERCOT record through Explanatory Comments (the white paper) and a Version Zero discussion draft shared at Large Load Working Group (LLWG). The submission did not treat NPRR1188 as background color. It treated the rule as the operating prerequisite and the missing planning link as the problem to solve. ERCOT had already determined that large CLRs required nodal dispatch and nodal settlement because their individual megawatt movements could materially change flows on specific transmission constraints. Version Zero asked planning to recognize that approved future Resource state. [3,4,26]

Version Zero closed the loop between nodal operations and large-load interconnection. An Interconnecting Large Load Entity could elect CLR study treatment, submit an upper operating level and a dispatchable floor, and attest that the future resource would register, qualify, follow SCED Base Points while online, and telemeter OUTL only when truly offline at zero consumption. The TSP would keep the full firm-service study and Load Commissioning Plan intact while testing whether dispatch toward the floor could mitigate binding N-1 constraints and permit earlier energization. [3,4]

The first draft used HSL and LSL because those were the familiar CLR operating parameters. The submitting load (interconnecting large load entity) would submit both values with the project information and attest to future NPRR1188 registration and qualification. In the steady-state study, the load would appear at its proposed output for N-0. For a binding N-1 contingency, the planner could redispatch the CLR toward LSL, including zero where elected. If the violation cleared, the study could identify the load as advanceable while the full-capacity firm study and Load Commissioning Plan continued. [4]

Version Zero also proposed a simple status discipline. The resource would be ON and subject to SCED whenever consuming. OUTL would mean what the word says: off-line and at zero. This removed an opt-out ambiguity that would have made the planning reliance fragile. Earlier energization had to correspond to a real operating obligation, not just when it was convenient to the load to respond.

The proposed language reached beyond one new definition. It added the CLR election to the large load information package, expanded Manual System Adjustment to include CLRs and ESRs, directed ERCOT to use a SCED tool when determining operating states for study purposes, allowed temporary interconnection configurations to be evaluated where reliable, and preserved the existing minimum-deliverability criteria for firm planning. The architecture was proposed in a format which has essentially stuck as the rules became formalized in new ERCOT Protocols: election, attestation, operating limits, study redispatch, continued firm planning, and private risk.[4] The economic case in the first PGRR134 draft argued that developers needed bankable study outcomes to unlock capital; TSPs needed model integrity and a way to focus regulated capital on durable regional upgrades; and customers should bear interim curtailment and compliance costs. This was not a request for free transmission, and this early clarity helped PGRR134 build allies with conventional industrial loads that were appropriately wary of new large loads with huge margins coming in to swoop up firm transmission capacity at preferential terms. It was a proposal to use available transmission more intensively while the project remained responsible for the consequences.[4,13]

PGRR134 Version Zero made simple

Study the full load for firm delivery. Permit earlier operation where N-0 is satisfied and binding N-1 thermal constraints can be mitigated by dispatching a registered nodal CLR toward its submitted floor. Keep building to firm.[3,4]

The accompanying white paper ("Explanatory Comments") were explicit about risk allocation. Loads electing CLR status would use system capability that otherwise sat idle, assume the risk of nodal SCED dispatch, remain in the traditional interconnection process, and receive firm service according to the Load Commissioning Plan. Private solutions could include gas turbines, storage, geothermal, dispatchable IT and balance-of-plant infrastructure, workload modulation, and software-enabled load shifting. Native load dispatch remained the response of last resort.[3,20]

These filings are the governing regulatory theory for what became PCLR in NPRR1325 and PGRR145. identifies the planning gap, grounds the solution in NPRR1188's nodal operating mechanics, preserves the Load Commissioning Plan, describes the private firming stack, places load dispatch last, and makes the project sponsor—not the ratepayer—the bearer of provisional delivery risk.[3,20]

PGRR134

Luminary filed the briskly refined concept as PGRR134 on November 1, 2025. ERCOT's public description captured the core architecture: TSPs would treat an NPRR1188-compliant CLR election as a study input and could authorize earlier energization when constraints could be mitigated by dispatch down to LPC, including zero where appropriate. PGRR134 was later withdrawn after ERCOT's Batch Zero package absorbed the concept into a broader system-wide process.[5,24,25]

PGRR134 established the freely scalable bargain that survived every later refinement: firm planning continues; the load's provisional operating band enters SCED; the customer assumes delivery risk; ERCOT retains the ability to dispatch consumption when the network requires relief. Any large load capable of qualifying and performing as a CLR can use the architecture, irrespective of the process or technology behind the fence line.[5,24,25]

PGRR134 also refined the language that would later become central to PCLR. Words we started with - HSL and LSL -- evolved into Maximum Power Consumption and Low Power Consumption. The filing protected earlier projects and existing large-load dates from impairment. It recognized that some reliability limits—including IROLs or other constraints not resolvable through CLR dispatch—would remain hard gates. It made clear that dispatchability could mitigate a defined class of operating constraint without claiming to cure every interconnection problem. [5,24,25]

The withdrawal of PGRR134 is itself a filing that documents institutional success. The sponsored proposal had done its work: it established the design space, demonstrated the minimum rule primitives, and gave ERCOT a tested solution to absorb. [5,25]

November 18, 2025: PGRR134 reaches PLWG

Agentic Infrastructure's companion Planning Load Working Group (PLWG) presentation placed the CLR election after steady-state analysis, CLR registration before the Quarterly Stability Assessment, and approval to energize after the Resource entered ERCOT's network operations model. It stated that the election could be applied retroactively to projects with completed steady-state analysis without restudy; that the CLR load would assume the full risk of interim curtailment; and that voltage or stability upgrades outside SCED would remain hard prerequisites. The deck also stated the design principle directly: the pathway was technology agnostic and could support front-of-meter or behind-the-meter co-location schemes. [24]

Luminary's sponsor presentation assigned the private developer two jobs: study its own POI exposure and assume the provisional delivery risk. ERCOT and the TSP retained the public study boundary. Thermal violations could be managed by advancing as a CLR at the ILLE's election; voltage and stability violations could not. Short-circuit requirements remained standard, and the Facility Study could adjust to the election. The CLR obligation lasted while the customer consumed above the firm floor and ended only after the firm upgrades in the Load Commissioning Plan were complete; the election did not change their scope or timing. [25]

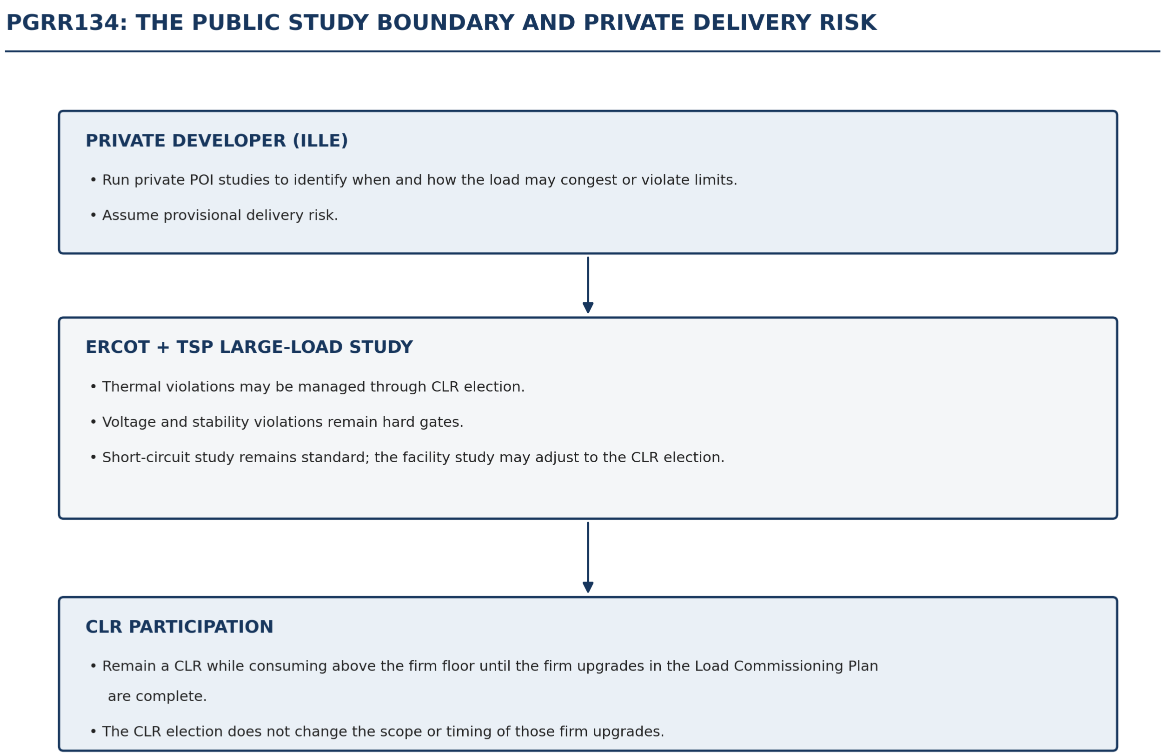

Figure 1. PGRR134 sponsor architecture: private POI diligence and delivery risk, utility study boundaries, and CLR duration through completion of firm Load Commissioning Plan upgrades. Adapted from Luminary Strategies, Nov. 18, 2025.[25]

Part IV — January to March 2026: Batch Zero Became the Vehicle

The January 26 interview

On January 26, 2026, the CLR architecture was discussed during ERCOT's Batch Zero interviews. This was the first interview that clearly tied PGRR134 into the batch conversations, eliciting that nodal CLRs were a useful operational solution for ERCOT rather than merely a worst-case curtailment tool. These resources could convert uncertain large-load growth into visible, controllable consumption; use operational headroom without displacing firm rights; and move congestion risk into a commercial decision by the project.

The distinction between flexibility and readiness also became clear. PCLR defines how a project will operate. Batch Zero eligibility determines whether the project is sufficiently mature to enter the transitional study. Site control, equipment commitments, utility work, tenant evidence, interconnection milestones, and financial security answer the anti-speculation question. A CLR election alone does not prove that the project is real, and that has proven true with the fundamental governance rules ERCOT has since enabled via protocol revisions in PGRR145 and NPRR1325, such as the new Form W.

Along the way in these interview processes, the CLR-in-exchange-for-speed bargain began winning in private conversations with developers. Survey results publicly available from ERCOT began to show rising numbers of supporters who confidentially pitched in favor of dispatchable solutions, even if those parties did not intend to invest in them directly.

The interview and survey phase of batch zero proceedings changed the frame inside ERCOT. CLR was no longer only the worst-case answer to a stressed grid and became a championed action item for ERCOT staff in grid operations, which is why it later became a feature of the batch zero rules. It was a positive planning and operating tool: use headroom when available, maintain a firm fallback, and convert curtailment into a Resource obligation rather than an emergency improvisation. That conceptual move is the origin of the word "provisional" which was added to the moniker, replacing the original "NPRR 1188-compliant CLR" language.

Batch studies enter the scene

The prior one-at-a-time study process had become unstable under simultaneous gigawatt-scale requests.[34] Each new load changed the background case for the next. Batch Zero emerged as a transitional system-wide study with five linked stages: inclusion, a common reliability study, a commitment gate, refinement around committed projects, and an actionable transmission plan. Readiness and financial commitment converted queue positions into demand the grid could actually plan around.[6]

PCLR gave Batch Zero a way to distinguish three quantities that an all-firm study would otherwise collapse: the load supportable with planning certainty now; the additional consumption that could operate conditionally; and the transmission expansion required for the full request. That distinction prevented the operational value of controllable load from disappearing inside an oversized firm-load assumption.

The Batch Zero commitment gate also changed the meaning of allocation. A planning model filled with uncommitted requests can force enormous upgrades and still fail to produce an actionable transmission plan. Requiring agreements, financial commitments, and acceptance of the ERCOT-defined allocation removes projects that are unwilling to stand behind their requests. The refinement study can then right-size the transmission solution around the remaining committed portfolio.[6]

That sequencing matters to PCLR because provisional access is valuable only if the grid can identify a credible path to firm service. The customer is accepting interim delivery risk, not purchasing a permanent exemption from transmission planning. Batch Zero pairs conditional operations with an eventual system plan. Connect, manage, and build occur in parallel sequence.

Senate Bill 6 and the fit for private underwriting

Senate Bill 6 rulemaking proceedings kicked off in yet another critical parallel sequence to the nodal controllable loads work. Within the legislation, core goals fulfilled by introducing in-market, price-based, voluntary dispatachability of loads included: support business development, minimize stranded infrastructure costs, and maintain system reliability. Luminary participated on multiple panels of the kickoff workshop at the Public Utility Commission of Texas on July 21, 2025. In those remarks, the firm's founder argued that Texas already possessed unusually sophisticated nodal operating tools but still lacked a planning-phase framework for allowing a large-load campus to take non-firm transmission access. She also discussed that approved Nodal Protocol reforms like NPRR1188 mattered only if they moved through implementation and into the engineering treatment of the future load. At the time, a battery paired with a new campus could still appear primarily as additional worst-case charging demand rather than as a controllable resource capable of reducing the campus's contribution to a transmission constraint.[35] The point was larger than giving large loads like data centers a legislated path to curtailment: rather, the point was to engage in reforms that reduced their curtailment exposure, allowing smaller system-wide obligations to govern over a large load class that responded favorably to preserve grid stability in real-time operations, arresting conditions that otherwise could lead to catastrophic outcomes for data centers and ordinary grid customers alike. She explained that the transmission study needed a credible way to recognize behavior that the operating system could actually enforce. In an analysis published two days after the workshop, she described the missing step directly: that Texas needed to credit controllable load and paired generation or storage in transmission studies, including studying batteries as offsets to firm load and developing netting arrangements that encouraged co-located resources to remain grid-visible. The governing principle was that curtailment capability should become planning-relevant where the real-time system could rely upon the corresponding behavior: in doing so, we allow voluntary regimes like connect-and-manage to grow, while turning large compliant loads in those regimes, into grid assets.[35]

PCLR ultimately reached that objective through a more scalable architecture than project-by-project attribution of a particular battery or generator. The load itself became ERCOT's enforceable Resource. The customer could then underwrite the provisional band with storage, generation, workload control, contracts, or other private infrastructure while ERCOT retained the ability to dispatch the load when those measures were insufficient. Senate Bill 6 therefore fit naturally with the bargain that emerged through PGRR134 and Batch Zero: the shared grid continues planning and building the transmission justified for firm service; the party seeking earlier access bears the incremental delivery risk above the grid-backed floor. [3,4,35]

Senate Bill 6 thus supplied the public-policy frame: support business development, minimize stranded infrastructure costs, and maintain system reliability. PCLR aligned those goals by keeping the customer in the firm interconnection process while placing the interim delivery risk above the firm floor with the party seeking earlier access. The customer could underwrite that risk through controls, storage, generation, contracts, and operating design. The shared grid continued building the infrastructure justified for the broader system. [3,4]

Part V — March to June 2026: ERCOT Builds the Minimum Sufficient Product

PGRR145 and NPRR1325

ERCOT filed PGRR145 and NPRR1325 on March 4, 2026. The package broadened the initial PGRR134 concept into the full Batch Zero process and added the operating rules ERCOT needed for PCLR: defined planning and operating quantity assessments which fit the new proposed batching process, commitment documents, stability treatment, SCED bidding, ramping, dynamic bid adjustment, and exit conditions. TAC recommended the package on May 19. The ERCOT Board recommended approval on June 2. The PUCT approved both revision requests on June 18. [1,7]

The final package consumed the governing architecture of controllable loads-for-speed at what can only be described as breakneck speed for power markets. At all times, it was clear that ERCOT could have independently turned NPRR1188 implementation into a robust market solution for CLRs that chose to be dispatchable: "connect-and-manage load" was always on the table. With Batch Zero adding layers of complexity, the effort across market stakeholders and ERCOT staff to fight for the inclusion of dispatachability in the final rules for Batch Zero can only be described as nothing short of wizardry.

Batch Zero thus became a system-wide allocation and transmission-planning process. PCLR became a binding provisional status for an otherwise eligible Studied Load within it. Under this construct, the customer/load resource could operate above its firm floor through SCED while remaining obligated to the study results, the interconnection agreement, the full stability case, and the ultimate transmission solution.

What the Final Package Carried Forward—and Added

Original filing	ERCOT final rules
Elective CLR treatment inside an individual LLIS	Transitional system-wide Batch Zero for otherwise eligible, mature and committed Studied Loads
MPC and LPC submitted with the election	MPC and LPC tied to Batch Zero allocation, registration and operation
N-0 at proposed output; N-1 redispatch toward the submitted floor	Firm allocation, full-MPC stability treatment, QSA/GTC controls and additional energization requirements
Attestation to NPRR1188 registration and SCED compliance, proposed exit obligations	Form W commitment, PCLR qualification, ramp-rate rules, bidding under "just in time last in line" SCED algorithm updates, and exit obligations
Private delivery risk and continued firm planning	Binding provisional status until the study-defined exit, with fallback to firm LPC
General planning reliance on CLR dispatch	Adjusted Bid Cap and specified SCED treatment to make congestion curtailment operationally enforceable [1,3,4,5,19]

The final rules thus placed the architecture inside a study, commitment, stability, registration, dispatch, and implementation system that ERCOT could administer across the entire backlog. Irrespective of whether the material rules of batch zero or its successor programs change or do not change, the ability to energize in advance of firm upgrades is a preserved contractual construct that can perpetuate in market design and remain accessible to any load that qualifies with compliance requirements.[1] Another important choice ERCOT accepted from the PGRR134 proposal, was that the planning and commitment provisions for electing PCLR as a status of future operation, could move forward while the PCLR-specific SCED (NPRR1188 system changes), Adjusted Bid Cap, ramp-rate, and energization provisions remained tied to later system implementation. On June 30 2026 , an ERCOT public market notice identified those deferred sections expressly, and thus PCLR has entered the grid code before every operating system change are complete, so that these commitments can be financed, made bankable for reliability planners, and self-evident in system planning exercises.[7,19] By July 11, 2026, Batch Zero could govern inclusion, study, commitment, allocation, and the transmission-planning process; remaining gray-boxed sections control the parts of PCLR that require changes to market systems. These include mitigated offer caps and floors, adjusted bid caps, specified SCED treatment, PCLR ramp-rate requirements, and additional energization and operating provisions.[19]

That fact should temper both criticism and celebration. PCLR has been adopted but it is not as of yet fully deployable in every operational detail. The market had crossed the institutional threshold while implementation work continues, which is a relieving trendline given how fractured the batch framework itself has become with recent political interventions. For a first-of-kind product, however, this sequencing is normal – the rules of the moat have to get written in stone so that building within it can rely on fundamentals, not politics.

The campaign for bankable flexibility to date has produced the minimum viable architecture ERCOT could approve on the Batch Zero timetable while preserving planning rigor and giving operations a real curtailment backstop. This is a presently enabled, future-state construct that creates global precedent for wires programs which seek to bring loads onto available grid headroom in exchange for enforceable load response.

Part VI — How PCLR Works

LPC and MPC

PCLR has two governing consumption levels. LPC represents the PCLR's grid-backed planning floor, bounded by the amount ERCOT determines can be reliably served for the applicable year. Maximum Power Consumption, or MPC, is the upper operating level requested and authorized for the Resource. Consumption between LPC and MPC is conditional. The PCLR submits an Energy Bid Curve and follows SCED Base Points within that band. The following diagrams provide a visual explanation of the firming band over the firm power limit assigned to a PCLR if the rules for Batch Zero are implemented as expected over the next eighteen months.[8]

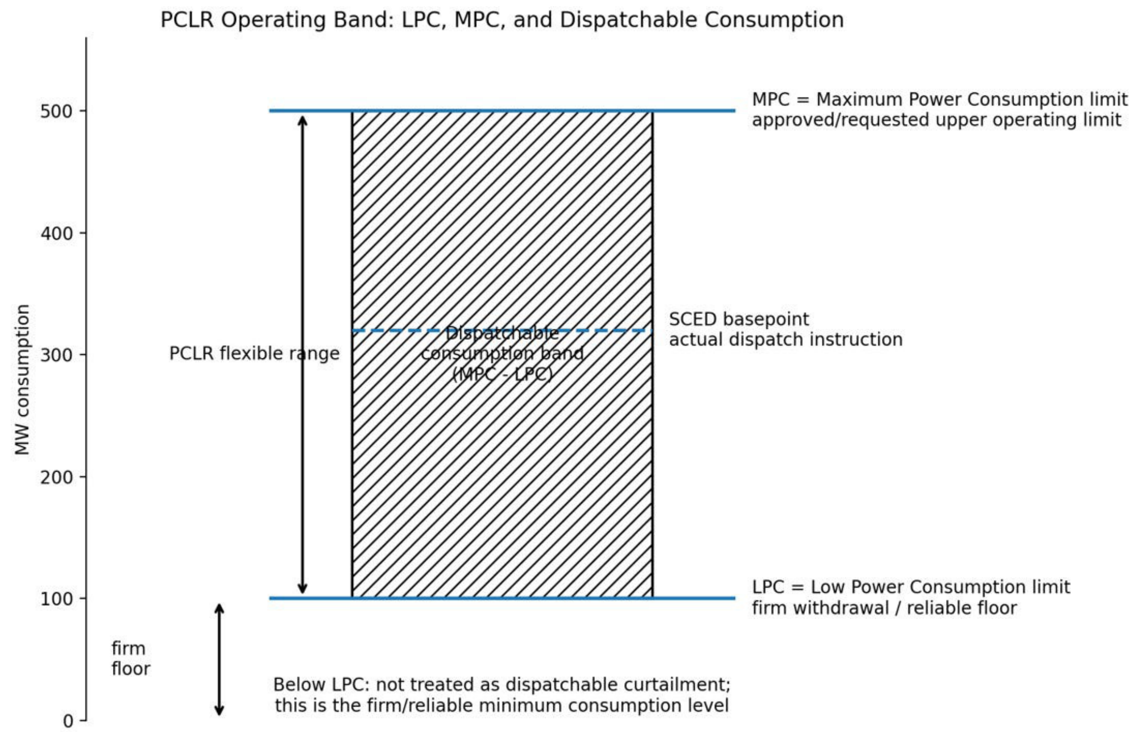

Figure 2. PCLR operating band: firm withdrawal up to LPC and dispatchable consumption between LPC and MPC. Source: Luminary Strategies, July 2026. [23]

Example: 500 MW requested; 150 MW firm

Requested load and MPC: 500 MW.

Batch Zero firm allocation and maximum LPC: 150 MW.

Conditional PCLR operating band: 350 MW between LPC and MPC.

The full 500 MW remains in the stability case and in the transmission plan for eventual firm service. [8]

Dynamic Bid Capping

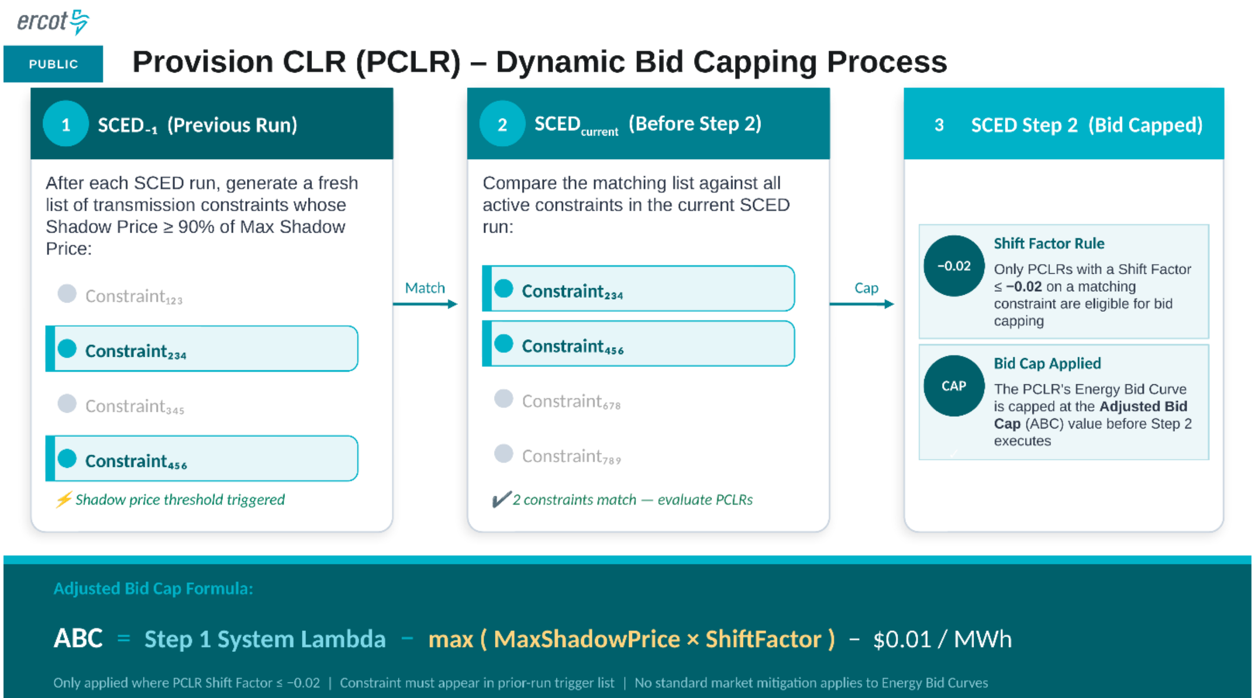

Figure 3: ERCOT's April 9, 2026 Batch Zero Workshop #7 deck, titled "Provisional CLR (PCLR) – Dynamic Bid Capping Process. See footnotes for link.

How to Read This Algorithm

ERCOT's dynamic bid-capping process is easiest to read as a three-run filter for deciding when a PCLR must become part of the congestion solution. The algorithm does not curtail a PCLR merely because a transmission element is heavily loaded. It first asks whether congestion has become sufficiently expensive and persistent to require the PCLR's bid to be modified.

1. **Look back one SCED run. In the previous SCED execution, ERCOT identifies transmission constraints whose shadow price reached at least 90% of the applicable maximum shadow price. This is the first trigger. Importantly, 90% refers to the shadow-price cap—not to 90% loading of a transmission element.**

2. **Confirm that the constraint persists. In the current SCED run, ERCOT checks whether a constraint identified in the prior run remains active. A transient constraint that disappears**

does not proceed through the PCLR bid-capping logic. The persistence test therefore prevents the special PCLR treatment from being triggered by a single isolated SCED result.

3. **Determine whether the PCLR can actually help. For a persisting constraint, ERCOT screens the PCLR's shift factor. A PCLR with a helping shift factor of ≤ −0.02 qualifies for the dynamic bid cap. ERCOT then calculates an Adjusted Bid Cap (ABC) and applies it to the PCLR's Energy Bid Curve before SCED Step Two:**

> ABC = Step 1 System Lambda − max(Max Shadow Price × Shift Factor) − $0.01/MWh

- **The modified bid does not itself dictate a particular curtailment quantity. It changes the PCLR's economics in the second SCED optimization so that the load is available to relieve the constraint rather than allowing its original willingness-to-pay bid to keep it consuming through congestion. SCED still determines the resulting Base Point through the full security-constrained optimization.**

- **In summary, ordinary economic redispatch gets the first opportunity to solve the constraint. The ABC mechanism becomes relevant when the constraint is expensive, persists across SCED runs, and the particular PCLR has sufficient locational sensitivity to relieve it. ERCOT therefore describes the purpose of the dynamic bid-capping process as ensuring that the PCLR is dispatched rather than allowing the applicable constraint to violate.[9]**

For each SCED execution, ERCOT has a PCLR's current market submissions and real-time operating data, including its Energy Bid Curve, Resource status, operating limits, ramp capability, and telemetry. ERCOT's Network Operations Model and State Estimator establish the current network topology and operating state. Network Security Analysis, informed by Real-Time Contingency Analysis, identifies and activates the transmission constraints and applicable limits used by SCED. SCED then evaluates the PCLR's bid-in consumption together with generation offers and other dispatchable Resources, subject to Resource and transmission constraints, and calculates a Base Point within the PCLR's currently feasible dispatch range.[2,12,26]

If no relevant constraint binds and the PCLR's bid supports continued consumption, the Base Point can remain at MPC. If congestion develops, ordinary redispatch may first move generation whose offers and shift factors make it economic to relieve the constraint. If the applicable PCLR constraint reaches the ABC trigger, persists into the next run, and passes the shift-factor screen, ERCOT modifies the PCLR bid for Step Two. SCED then determines the lower Base Point required by the full optimization.[9]

The Qualified Scheduling Entity (entity owning technical interface between the load site controller software and grid operator software) receives the Base Point and the load follows it within the applicable ramp requirement. The customer's private controller may discharge a battery, start generation, transfer a load block, defer workload, or reduce native consumption. ERCOT does not need to choose the internal asset: the idea is to remain agnostic to what is "behind the PCLR". The grid operator only needs the registered PCLR to reach the instructed consumption level and continue providing accurate telemetry.

The next interval starts with the changed system condition. The PCLR may remain curtailed, move lower, hold its level, or begin a SCED-authorized restoration. This repeated loop is the operational substance of conditional access. The customer does not decide that transmission headroom exists and self-dispatch upward. It offers availability to consume. ERCOT authorizes the quantity through the Base Point.

The operating loop of SCED is no easy feat: achieving it, however, means that the project availing of grid headroom around peak can do so freely, irrespective of bad news received from the batch zero process. PCLR is largely agnostic to the eventual firm-allocation outcome. Batch Zero allocation governs the grid-backed floor; it does not govern how much of the approved load ramp may be available for conditional consumption if PCLR election has been made on Form W. ERCOT determines the amount it can reliably serve as firm load and that amount binds the firm power limit, or, LPC. MPC, by contrast, can remain tied to the PCLR's requested load ramp for the applicable year. Thus, if a 500 MW Studied Load receives only 150 MW of firm allocation, the remaining 350 MW does not disappear from the operating architecture. It remains the provisional band between LPC and MPC, available for consumption when the Resource's Base Point and prevailing system conditions permit. The same architecture works whether the customer receives a large firm allocation, a small one, or sees its firm allocation change materially from year to year. A qualifying Batch Zero Studied Load can, subject to its TSP interconnection agreement, be configured as a PCLR for dispatch up to the requested load ramp—MPC—for that year even where ERCOT cannot allocate the full request as firm service. If ERCOT cannot identify a timely transmission solution for some portion of the request through the Batch Zero horizon, that unserved portion may remain unallocated for firm service and be considered in a later batch; the absence of allocation applies to LPC, not MPC.

What limits consumption in that delta is instead the operating system described above: local congestion can cause SCED to dispatch the PCLR downward, and stability limitations identified through the Quarterly Stability Assessment can be enforced through GTCs in SCED. System-wide congestion risk also has to be accepted, which means a site must study to optimize the potential upside of PCLR dispatch over the downside of waiting years for power. In practical terms, allocation determines what the grid will promise; PCLR is a voluntary election to be dispatchable when private parties determine the location is favorable to a grid utilization thesis that enables: (i) consumption of a higher amount of power consistently, and (ii) self-managing solutions when congestion binds.

The PCLR election is an engineering program shaping risk and reward

PCLR is not a demand-response program selected for a few hot afternoons. The facility becomes a dispatchable Resource. It needs a QSE, real-time telemetry, dispatch-following controls, truthful operating limits, and a power-management system capable of translating a five-minute Base Point into an orderly sequence of workload reduction, storage response, generation, switching, and restoration. [1,22]

The models must show the behavior ERCOT is relying upon. Where applicable, PSS®E and PSCAD representations should capture controllable-load response, ramping, disturbance behavior, and recovery rather than leaving the facility as a constant-load block. Ride-through and dispatch obligations coexist under the broader ERCOT large-load reliability framework. The engineering representation of the facility therefore has to remain consistent with the operating behavior ERCOT expects from the Resource. [22]

Form W, now available in ERCOT Nodal Protocols § 23, makes the election staged and binding: the project first declares its intent and minimum LPC by study year, then accepts the study's resulting LPC and Exit Date with the interconnection agreement. Until the Exit Date, the provisional band remains subject to the PCLR operating restrictions, including the restriction on providing Ancillary Services. [1,7,22] ERCOT published Form W as part of the current Nodal Protocols effective July 11, 2026.

A key power-engineering feature of PCLR mirrors the logic ERCOT has long applied to generators seeking speed to power through connect and manage: an asset may interconnect without receiving a guarantee that the transmission network can accommodate its maximum capability in every hour. The difference is direction. A generator accepts the risk that some otherwise available output cannot be injected; a PCLR accepts the risk that some otherwise desired consumption above LPC cannot be withdrawn. In both cases, physical interconnection and full firm deliverability are separated, while congestion and curtailment remain operational consequences of the network.

This is the critical consequence of separating firm allocation from operational capability. A conservative planning result can remain conservative without forcing every unallocated megawatt to remain dark in every unconstrained hour when the transmission system can physically support additional transfers. The customer receives no promise that the provisional band will be bankable twenty-four hours a day, just as a connect-and-manage generator receives no promise that every available megawatt can always reach the market. In each case, commercial performance improves as usable transmission capability expands and congestion exposure falls. ERCOT's generation-side experience demonstrates why transmission expansion and connect-and-manage belong together: operational congestion can accelerate market access, while persistent congestion also reveals where additional transmission creates value.

That relationship clears up a common misconception about dispatchability. Dispatchable resources are not substitutes for more wires or greater transfer capability. They are how operators use a network whose physical capability is finite and changes with topology, outages, generation patterns, and contingencies. More transmission increases the amount of power the system can move; dispatchability determines which injections and withdrawals should move when a particular limit binds. Even a larger and stronger network must still survive G-1 and N-1 conditions, and SCED still needs Resources whose locational response can relieve the constraint that actually appears.

The two capabilities – more dispatchable assets and more wires and wires transfer capability – compound: more wires create more usable headroom, while better dispatch allows ERCOT to use that headroom safely and economically to maximize grid utilization. More transmission increases the amount of power the system can move; dispatchability determines which injections and withdrawals should move when a particular physics limit binds the wires system.

The response is nodal and network-wide

A PCLR"s measurable consumption enters SCED at a Resource Node and changes flows across monitored transmission elements according to nodal shift factors. A Base Point reduction at one load so participating at a node can relieve a local line, a broader network constraint, or a GTC elsewhere in the study area when the modeled sensitivity is material. The POI remains the electrical and metering boundary through which the site's net withdrawal is observed and the response at the POI is reactive to the known elements of the transmission network model.[2,3]

This is why NPRR1188 was indispensable. Gigawatt-scale load movement cannot be managed through a generic zonal assumption. Nodal dispatch and nodal settlement make the consumption visible where it actually affects congestion.

The distinction between local and system-wide congestion is important for private firming. A co-located battery can have a strong one-for-one effect on the site's grid withdrawal and on nearby thermal loading. Its effect on a distant network constraint depends on the same shift-factor physics governing every Resource. PCLR does not promise that onsite supply neutralizes all system constraints. It ensures that the load itself remains available to ERCOT when the private assets co-located and separately metered from the load do not deliver the needed network relief.

The Adjusted Bid Cap

A high PCLR bid expresses a high value of continued consumption. It does not guarantee that SCED will reduce the load when a transmission constraint reaches its modeled limit. Maximum shadow-price caps can cause the optimization to tolerate or relax a constraint rather than dispatch an expensive bid-in load. ERCOT therefore designed the Adjusted Bid Cap, or ABC, as the operational bridge between the PCLR planning assumption and real-time curtailment risk.[9]

How ABC enters SCED

1. ERCOT looks back one SCED interval and identifies constraints whose shadow prices reached the specified percentage of their maximum shadow-price caps.

2. Only constraints that remain active in the current run enter the match list.

3. ERCOT screens whether the PCLR has a qualifying helping shift factor. The workshop draft used 2%.

4. ERCOT calculates the applicable cap across matched constraints, selects the minimum result, and subtracts one cent from the Step-One system-lambda-based value.

5. The modified bid enters SCED Step Two. SCED still determines the Base Point.[9]

ABC is therefore a reliability override on the submitted bid, not a forecast of ordinary curtailment. A PCLR may bid the highest permitted value across its operating band because the customer wants to consume whenever the system can support it. The cap intervenes only for the qualifying constraint. The customer's willingness to pay remains economically meaningful everywhere else.[9]

Multiple constraints complicate the result. A PCLR may materially help more than one active constraint, each with a different maximum shadow-price cap, shift factor, and Step-One condition. ERCOT described taking the minimum applicable cap. The most restrictive qualifying constraint therefore governs the modified bid. This is operationally conservative and commercially difficult to model because the relevant constraint set changes with topology, outages, generation dispatch, and later interconnections.[9] The final diligence model must therefore look beyond the transmission upgrades identified in the load study and examine the broader SCED constraint footprint.

The 90% trigger refers to 90% of the constraint's maximum shadow-price cap: ABC responds to the economics of the modeled constraint, and maximum shadow-price caps vary by voltage level. A lower-voltage constraint may therefore activate the mechanism earlier than a developer expecting exposure only to a major regional contingency.[9]

The shift-factor screen also makes PCLR curtailment risk inherently location- and constraint-specific. Every Resource has a calculated shift factor on each modeled transmission constraint, but only material sensitivities should affect the PCLR's dynamic bid cap. Those same sensitivities can also inform how the customer sizes its private insurance against curtailment. A load may choose to install more onsite firming, storage, generation, or workload-shifting capability than would be required to cover the expected worst curtailment event at its immediate location, particularly where modeling shows that constraints elsewhere on the ERCOT system could also materially affect its dispatch. In that sense, the relevant engineering question is: "[a]gainst which modeled constraints is this Resource materially exposed, how often might those constraints bind, and how much response does the customer want available when they do?"

Geography alone therefore does not define the risk. A load in Midland may have a nonzero shift factor on a Houston-area constraint, but if that sensitivity is too small to pass the applicable screen, that constraint should not drive PCLR bid capping or the sizing of a mitigation portfolio. Conversely, a geographically distant constraint can matter if the Resource has sufficient electrical sensitivity to it. The useful risk map is electrical, not merely geographic: expected binding constraints, their severity and frequency, the Resource's shift factor on each, and the customer's desired protection against the resulting curtailment exposure together determine how much private firming or load-shifting capability is worth buying.

Example

Why load can curtail before a willing generator

Generator offer: $800/MWh; relieving shift-factor magnitude: 10%.

Implied shadow price to obtain the relief: $800 ÷ 0.10 = $8,000/MW.

Assumed maximum shadow-price cap: $3,000/MW.

PCLR-supported shift-factor magnitude: 2% (ERCOT screen: ≤ −0.02); shift-factor-weighted congestion component: $3,000 × 0.02 = $60/MWh.

At this offer and sensitivity, re-dispatching the generator would imply $8,000/MW of constraint relief—well above the assumed $3,000/MW maximum shadow-price cap. If the constraint reaches the ABC trigger and the PCLR passes the helping shift-factor screen, ERCOT may cap the PCLR's Energy Bid Curve and SCED may curtail the load before re-dispatching that generator. [15]

This example exposes the limit of the phrase "last in line." The current ABC is a constraint-triggered bid-mitigation mechanism. It approximates the policy goal of preserving PCLR consumption until congestion truly requires the load. It does not comprehensively compare the PCLR against every generation offer and operating limit that might relieve the same constraint.

"Just in time" is more concrete. The constraint must be active, highly priced, persistent into the next interval, and materially affected by the PCLR.

In the medium-term, aspirational improvements to the PCLR architecture should compare resource offers, shift factors, high-sustained limits, constraint severity, and private mitigation associated with the load before reducing PCLR consumption. The gap becomes acute when the customer has built private mitigation that simply does not account for this calculus. Suppose a battery is offered into SCED at a price intended to make it move before the PCLR curtails. Once ABC lowers the load bid, the merit order can invert: native load reduction becomes cheaper than using the full battery response. ERCOT obtains reliable relief, but the customer loses the operating sequence it financed. That is not a reason to discard ABC, but it is indeed reason to refine the market design carefully and intentionally as system changes are scheduled to implement NPRR1188 and NPRR1325.

Minimum sufficient, not final

ABC gives ERCOT confidence that PCLR load will move when a qualifying constraint requires relief.

A complete last-in-line design would also preserve the economic priority of verified private mitigation and compare all relevant redispatch options. These ideas are discussed further below.

Part VII — Co-Location and the Private Firming Architecture

The co-location myth

One of the earliest Batch Zero implementation problems was semantic. ERCOT workshop slides described "Load-Only CLR (No Generation or Batteries)" and CLRs "not netted with any Generation or Batteries." Developers heard a site-wide prohibition on co-located storage and generation. The actual technical boundary was narrower.

A battery or generator cannot be the mechanism by which the load qualifies as a CLR, and those resources cannot be commonly metered with the CLR in a way that obscures CLR compliance. A separately registered, separately metered, and separately settled Resource can still sit behind the same high-voltage POI. Luminary asked ERCOT to make that distinction explicit. ERCOT staff agreed that the clarification made sense.

The result separates compliance architecture from electrical architecture. The PCLR must satisfy its own registration, telemetry, dispatch, metering, settlement, and performance duties. The customer may coordinate batteries, generation, segmented load, buses, transformers, feeders, and switching schemes behind that boundary. Separate market visibility can coexist with coordinated physical operation.

This separation is not a drafting technicality. Common metering can obscure whether the consuming Resource itself followed the Base Point. Independent metering preserves compliance evidence. A common high-voltage POI can still support electrical coordination below that boundary. The site can share switchgear, buses, protection zones, transformers, and controls without collapsing unlike market Resources into one settlement object.

This results in a rather practical division of labor: ERCOT dispatches the registered Resource and the customer orchestrates the site, making sensible decisions to "cover" for potential load curtailment by oversizing paired assets or oversizing the private "firming" that protects key commercial workloads from ever experiencing disruptions. Redundant transformers, sectionalized medium-voltage buses, and normally open ties – standards in intelligent campus designs pairing loads and generators, have tremendous relevance to the private firming vision enabled by PCLR. Switching scheme usually built for last-minute emergency power solutions will become commonplace to assume firming duty when loads are called on for everyday grid support. In prior articles and lectures, Luminary has dubbed this "Bring Your Own Everything." The implication is that energy storage responds first and reduces net grid withdrawal while preserving customer operations. Onsite generation extends the coverage when battery duration runs out. Dispatchable load becomes the final layer when the constraint lasts longer than the private firming stack. Equipment redundancy and resource diversity confine a failure to a portion of the facility rather than converting one asset outage into a full-campus interruption.

This is the correct engineering description of the product: preserving headroom consumption while protecting the POI withdrawal limit. The site does not hide the gross load or eliminate the registered PCLR. It uses private infrastructure to keep the measured grid-facing behavior inside the operating envelope ERCOT has authorized.

A modular design also makes the commitment more credible. A single battery sized to the full provisional band creates a single duration and outage dependency. Multiple trains permit the site to declare only the capability actually available after equipment outages, state-of-charge limits, maintenance, or fuel restrictions. Redundancy converts an all-or-nothing flexibility claim into a measurable operating envelope.

The hierarchy is economic as well as physical. Battery discharge is used first where preserving compute is worth more than cycling cost. Generation follows where duration, emissions, fuel, permits, and operating cost support it. Workload modulation enters where the Value of Compute Load is lower or the service contract permits delay. Native load reduction closes the gap. The site can change the internal ordering over time while the PCLR obligation remains stable.

The minimum sufficient PCLR does not yet guarantee that a privately procured battery will protect the associated load in SCED merit order. Market design enhancements can support that outcome now that the base PCLR design has been written. Today, ERCOT would the separately metered battery and PCLR as separate Resources. SCED may dispatch the battery, recognize its relief across the network, and still curtail the PCLR because the battery's contribution is not attributed one-for-one to that load. The project pays for the asset while the benefit is socialized across the footprint. A developer must be able to finance a battery as a dependable curtailment hedge; ABC implementation if does as articulated today by ERCOT, operates such that the PCLR bid could be low enough to move native load reduction ahead of the battery.

Curtailment risk axis

In a recent research partnership between Luminary and Aurora Energy Research, it has been discussed that a SCED assessment can usefully separates PCLR curtailment into two distinct drivers: local transmission congestion and system-wide scarcity. That distinction matters for private firming because a resource can be highly effective against one and largely irrelevant to the other. Local congestion is a locational problem. System-wide scarcity is an adequacy and dispatch problem. A PCLR can face either condition independently, or both at once.[32]

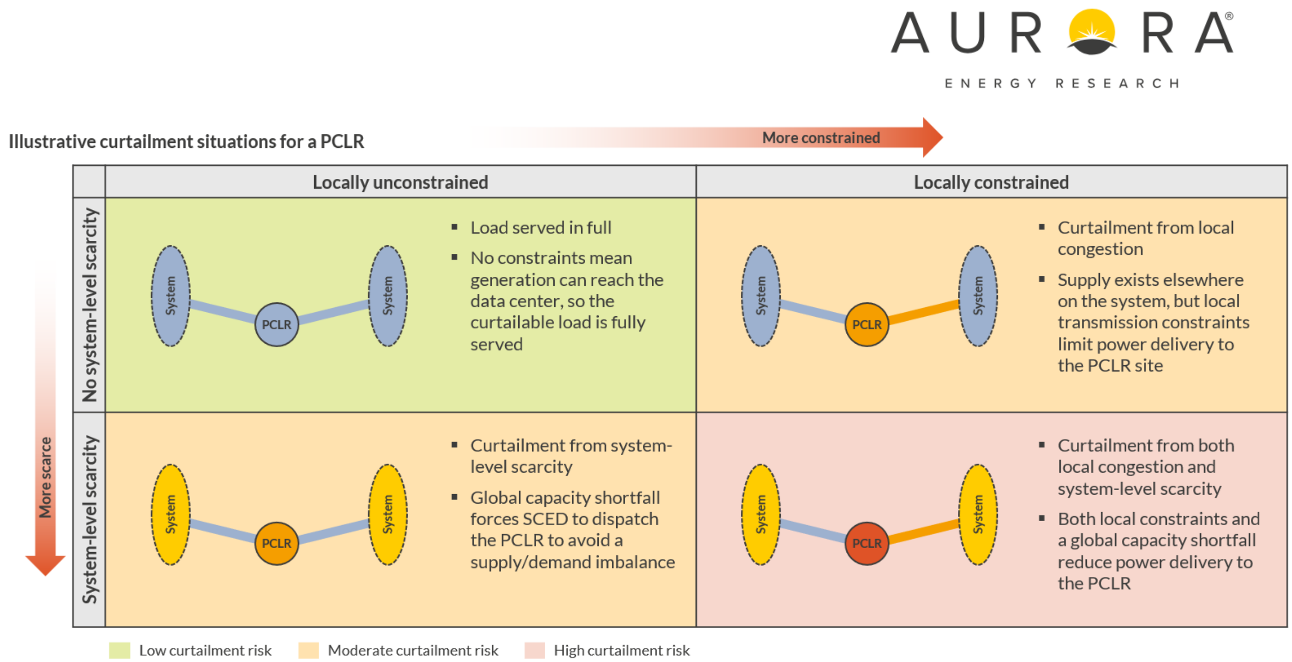

Figure 3. Illustrative PCLR curtailment regimes under combinations of local transmission congestion and system-wide scarcity. Replicated with permission from Aurora Energy Research, "*SCED Model Assessment: How Local Transmission Constraints and System-Wide Scarcity Can Drive PCLR Curtailment*," provided Aug. 10, 2026; (forthcoming publication expected Aug. 26, 2026).

The distinction also bounds what attribution can accomplish. A battery or generator with the right locational effect can preserve headroom against local congestion if its mitigation is actually attributed to the sponsoring PCLR. It cannot create system capacity that does not exist, and it cannot guarantee relief from a system-wide scarcity condition it does not solve. The next market-design step is therefore not generic "firming." It is constraint-specific attribution: preserving the financeable value of private mitigation where the asset physically relieves the condition driving PCLR curtailment, while leaving residual system exposure with the customer. [32]

Front-of-meter attribution as the next design step

Illustrative one-for-one attribution

Total load: 200 MW.

Batch Zero planning LPC: 100 MW.

Co-located, separately metered battery: 100 MW.

When the battery is demonstrably operating and delivering the relevant offset, the QSE could establish a temporary effective dispatch floor of 200 MW, supported one-for-one by telemetry. The planning LPC remains 100 MW.

This front-of-meter netting or attribution concept remains future-state work. It does not permanently increase the Batch Zero planning LPC or require ERCOT to accept a private contract as transmission service. It asks ERCOT to recognize verified physical mitigation through a higher temporary dispatch floor while the paired Resource is actively performing. When that support disappears, the floor falls and the unsupported portion of the load returns to ordinary PCLR exposure. The principle is straightforward: private mitigation should retain enough attributable value to be financeable. [23]

The design question is evidentiary. ERCOT should not increase the firm floor merely because the customer has a contract with a generator. It can recognize a temporary operational floor when telemetry proves that a separately metered Resource is actively offsetting the load in the amount and direction required by the constraint. Ex-post validation, penalties for unsupported adjustments, and clear QSE responsibility can preserve reliability while preventing the economic value of private mitigation from evaporating into a socialized system benefit.

This is where front-of-meter netting differs from the original settlement question that started the journey. The earlier question asked whether netting could change planning treatment (the answer here was “no”). The latter proposal asks whether verified operating relief can be attributed to the associated PCLR inside SCED. The latter may be essential to making speed-to-power products financeable.

Part VIII — Dispatch Inside Commitment

Flexible does not automatically mean dispatchable

A large load can possess technical flexibility and still fail as a grid Resource. The operating service must specify all of the parts of what is expected of CLRs today in ERCOT, and what will be expected of voluntary election-enabled PCLRs: ramping and curtailment response magnitude (difference between MPC and LPC), response time, ramp rate, duration, notice, minimum time between events, recovery, rebound, dependencies, telemetry, verification, and fallback behavior. The same facility may be useful for day-ahead shaping, unhelpful for a fast contingency, and valuable again for a prolonged event when paired with generation. [10]

The familiar distinction between unit commitment and economic dispatch helps frame the burden. Commitment establishes which capability is available across time and under what time-coupled conditions. Dispatch determines how much of the available capability should be used in the interval. For PCLR, the commitment layer lives across Resource status, operating limits, QSE submissions, workload availability, battery state of charge, generator readiness, and private firming contracts. SCED dispatches inside that credible envelope. This is why "[t]elemetry and verification convert capability into evidence." [10]

Recovery is part of deliverability

Bankable dispatch requires a multi-interval view of availability, duration, and restoration. A five-minute reduction can create a later problem. Deferred compute must run. Batteries must recharge. Cooling and thermal storage must recover. Repeated events can shrink the next interval's capability. A response that solves the present interval may move the peak unless recovery is represented.[10]

Restoration is particularly important for compute. A campus may be physically able to drop 300 MW in seconds by shedding a block, but incapable of restoring the same block at the identical speed without disrupting workload, cooling, or supply balance – which is why our formulation respects the work of "Bring Your Own Everything" to include computational flexibility on the bottom end of dispatch, so that it has time to prepare to respond and recover after native generator and battery curtailment options. Treating upward and downward movement as symmetrical can create false performance expectations and repeated Base Point deviations.

The same principle applies to batteries. Fast response says little about sustained response. State of charge, degradation policy, reserve obligations, and recharge windows determine whether the capability remains available in the next hour. Commitment establishes the time-coupled envelope. SCED decides how much of that envelope to use now. [10]

The 500/100/470 MW operating example

Dynamic MPC after curtailment

Initial state: LPC = 100 MW; MPC = 500 MW; Base Point = 500 MW.

Congestion event: SCED issues a 470 MW Base Point; the PCLR reduces to 470 MW.

Slow restoration case: the PCLR resets MPC to 470 MW. The next run can hold or reduce the load, but cannot immediately send it back above 470 MW.

When the constraint clears, the PCLR raises offered MPC gradually and waits for SCED to authorize higher Base Points. It does not self-dispatch upward.

This example addressed the anti-oscillation problem created by symmetric CLR performance rules. A large load may shed rapidly and restore slowly. Dynamic MPC prevents SCED from instructing an immediate rebound above the level the facility can physically restore. The operating obligation remains strict on downward dispatch. Upward consumption returns only through SCED authorization.

A more mature PCLR design should formalize that asymmetry. Fast curtailment protects the grid. Controlled restoration protects frequency and avoids performance penalties for a physical capability the resource never claimed. The rule should reward truthful operating limits rather than fictional symmetry.

Part IX — The Intellectual Architecture Behind PCLR

In 1978, Fred Schweppe described a power system in which customers, utilities, storage, generation, communications, mathematical models, spot prices, and physical control formed one hierarchical energy marketplace. System-level controllers established targets. Customer-side controllers decided how to buy, sell, store, generate, and consume within those targets. Prices provided soft load control; interruptible rights provided hard control when system security required it.[11]

Billimoria and Byers place work by Luminary and others in the modern Schweppian lineage: demand bidding its flexibility and being dispatched through SCUC and SCED directly or through intermediaries. That placement is precise. PCLR converts large-load flexibility from an administrative promise into bid-in, nodal, security-constrained demand.[15] The authors also describe a Value of Compute Load (VoCL) as the number that should sit at the core of a non-firm interconnection or flexibility agreement because it determines both the opportunity cost of entering the agreement and the real-time incentive to comply with operator instructions.[16]

This paper agrees that PCLR is a modern expression of that architecture and that VoCL is the economic bargain that is made possible by PCLR and similar constructs. ERCOT's control room identifies network conditions and issues the Base Point. The QSE and site controller translate that instruction into operating limits and private Resource actions. Batteries, generators, breakers, workload systems, and cooling controls execute the physical response. Measurements travel upward. Targets travel downward. The customer becomes a controllable economic actor inside secure system operations.

PCLR extends that architecture to the present load-growth problem. Schweppe's customer was already connected. PCLR uses the same control architecture to govern access before the customer's full firm transmission allocation is available.

Byers and Billimoria discuss the utility system advancement implications of Schweppe's work in their recent publications. The work notes that demand capacity is an investment variable. Large industrial customers choose scale, location, reliability, risk appetite, contracts, onsite supply, and flexibility. Open access remains bounded by security-constrained dispatch for both generation and load. Efficient prices coordinate those decisions; hedging allocates the volatility to the parties able to bear it.[13]

PCLR certainly operationalizes that aspirational framework. The customer receives provisional access subject to nodal SCED just like generators do today in ERCOT. The firm interconnection process continues building toward full delivery, which appropriately distinguishes the economic incentive of load to consume from the generator incentive to maximize economics. The sophisticated large-load customer bears the price and congestion exposure above LPC and chooses its own hedge: storage, generation, supply contracts, workload controls, an energy service agreement, or some combination. Existing consumers do not automatically underwrite the customer's private upside.

The Harvard policy brief (Billimoria and Conleigh, March 2026) is thus especially important where it treats demand capacity itself as endogenous, which is precisely how firming choices for PCLR will work. The large-load customer chooses whether to build, where to build, how much to build, how much reliability to buy, how much price exposure to retain, and which hedge to procure. PCLR makes those choices legible to the power system. LPC represents the PCLR's grid-backed planning floor, bounded by the amount ERCOT determines can be reliably served for the applicable year. The provisional band is the amount of operational access the customer is willing to underwrite under SCED.[13] Billimoria and Byers correctly place security-constrained dispatch at the boundary for both generation and load. That principle is the market-design heart of PCLR. Access can be open because the system retains authority over the quantity that actually clears. Network physics—not a contractual fiction of full deliverability—governs the operating outcome.

The Value of Compute Load (VoCL) concept from the authors further gives economic context to the PCLR bid. Billimoria and Byers translate observable cloud prices into the short-run marginal benefit of electricity consumption for compute. Their estimates vary sharply by hardware vintage, service quality, location, time, workload, utilization, and SLA exposure. Spot compute, on-demand compute, and SLA-backed services do not carry one common value.[14]

That heterogeneity is exactly what bid-in demand is designed to reveal. The Energy Bid Curve can express which consumption is valuable enough to remain online, which workload can yield to congestion, how service priority changes the answer, and when private firming is cheaper than native load reduction. A generic Value of Lost Load or static flexibility tier cannot do the job.

The marginality discipline expressed in their work matters as well. Ordinary amortized data-center capital does not belong in a short-run bid. A real intertemporal opportunity cost may belong where present operation consumes future value: battery degradation, scarce cycle life, accelerated hardware wear, thermal recovery, or deferred-work obligations. PCLR fits into this thinking because the design solution creates an incentive to build hardware and software capable of repeated, measurable dispatch that acts as an insurance policy over and above what the site analysts expect the load may experience in grid-triggered curtailment.

The numerical spread in the VoCL work confirms why a single administrative flexibility value is inadequate. Interruptible and spot compute can carry materially lower short-run value than on-demand service. Mainstream and frontier accelerators can differ sharply. SLA exposure can push the economic cost of interruption far above ordinary power-market caps. Location, utilization, workload, and service priority change the answer again. The PCLR bid is valuable because it can represent this heterogeneity where the dispatch decision is actually made.[14]

VoCL also changes the meaning of private firming. A battery is economic when the cost of discharge and foregone future use is lower than the compute value preserved. A generator is economic when fuel, emissions, maintenance, and availability costs are lower than the value of continued load. Workload modulation is economic where deferral costs less than either physical hedge.

Physical and Financial Hedging

PCLR creates two linked hedge markets. Physical hedges change the site's electrical behavior: batteries, generation, redundant feeders, switching, thermal storage, and controllable workload. Financial and contractual hedges change who bears the economic consequence: power contracts, availability guarantees, service-level agreements, energy service contracts, options, QSE portfolio arrangements, and insurance-like products around curtailment. [13,14]

The spot signal remains essential even where the customer seeks certainty. A hedge without a credible underlying price or dispatch exposure becomes an administrative subsidy. A hedge around nodal SCED allows the customer to buy the amount and shape of certainty it values while leaving the residual operating quantity visible to the market. That is how private underwriting can protect existing consumers without muting the investment signal for new generation and transmission.

The best PCLR commercial products will translate variable grid access into a defined level of customer service through technology, controls, and contracts. The moat belongs to firms which can make that translation credible to ERCOT, the TSP, the QSE, the tenant, and the capital provider.

Administrative curtailment leaves the market out

Administrative programs begin with a presumed baseline and impose a curtailment right through classifications, contracts, emergency triggers, or operator discretion. They must solve baseline integrity, incentive compatibility, compliance, and rationing outside the ordinary dispatch process.

This difference matters for PJM and for vertically integrated utilities now reaching for a "flexibility" tariff. Markets and utilities with security-constrained commitment and dispatch capability have the essential machinery. PCLR sets a new bar for clearing supply and bid-in demand together through transmission-security-constrained economic dispatch, supported by the commitment processes that determine which capabilities are available to be dispatched. The grid should not wait until emergency shed to discover whether the capability is real. [2,10,13]

The administrative approach is tempting because it appears simpler. The utility writes a right to curtail into a tariff and treats the obligation as available. The simplicity disappears in operations. Which megawatts move first? What baseline proves the reduction? Does the load have the duration? Can it restore without rebound? Is the operator instruction economically aligned with the customer's immediate interest? What penalty is large enough to overcome the value of continuing compute? A market-integrated Resource answers these questions through bids, limits, telemetry, Base Points, and performance rules.

Administrative rights will remain necessary for emergency reliability and firm load shed. They should not be the primary architecture for ordinary congestion management when the same load can participate in the normal dispatch continuum. The goal is to use emergency authority less often because the market has already found the least-cost, physically valid response.

Part X — Grid 2.0: From Administrative Curtailment to Embedded Dispatch

PCLR raises an obvious implementation question. If large-load curtailment is going to migrate over time from bespoke utility agreements, emergency calls, interruptible tariffs, and other administrative interventions into ordinary grid operations, what technical architecture lets thousands of flexible megawatts actually respond in a repeatable way?

Contemporaneously with the development of PCLR, the Grid 2.0 Task Force began working on that problem. Its first Request for Comments (RFC), published in August 2026.[21] Grid 2.0 is an open protocol standard for orchestrating grid access among flexible loads, storage, supply, and transmission constraints. Its authors explicitly identify ERCOT PCLR as one of the emerging market structures the protocol is intended to complement. In fact, the RFC says that initial Grid 2 deployments are intended to serve as an "operational layer to realize the verifiable dispatch capabilities expected in ERCOT PCLR." The distinction is useful. PCLR establishes the grid-facing market obligation: the Resource bids to consume, ERCOT determines the Base Point, and the PCLR must follow it. Grid 2.0 asks how the customer might organize everything underneath that obligation so that the required response occurs quickly, automatically, and predictably.

Its answer is to insert a voluntary orchestration layer between the grid operator's slower market-dispatch interval and the facility's faster controls. Grid 2 operates on a proposed one-minute clock. In Luminary's experience, the idea of an intermediate synchronized time clock is a missing piece of the dispatch puzzle that operators of technology behind the site controller are indeed struggling to determine.

In the Grid 2.0 proposal, flexible load appears as a Use resource; storage becomes a Buffer; generation becomes a Source; and an Allocator distributes permission to use available grid headroom among participating resources. Loads can divide consumption into differentiated service classes, while storage is designed to clear before native load must be curtailed. If the orchestration system loses its signals or context, the participating load falls back to its existing firm or Grid 1 baseline.

That is where Grid 2.0 becomes particularly relevant to PCLR. PCLR replaces an administrative question—*whom should we call and ask to curtail?*—with an operational one—*what consumption quantity can this Resource support under the network condition that exists now?* Grid 2.0 begins sketching the next layer: once ERCOT has established the quantity available to the PCLR, software can determine how that quantity is implemented across compute workloads, batteries, onsite generation, and other controllable equipment. The RFC deliberately places this one-minute orchestration below typical five-minute market clearing and above sub-second grid controls, leaving the existing grid operator, market, and reliability architecture intact.

The two ideas therefore fit together: such that a dispatchable-load market object sits on top of what Grid 2.0 posits as a common technical language for orchestrating the assets that make dispatchability bankable. The RFC illustrates how the industry is already thinking beyond one-off administrative curtailment toward flexible load as an embedded operating capability: continuously visible, automatically allocated, buffered where possible, and ultimately responsive to the same physical constraints that govern the rest of the power system.

Part XII — The Moat and the Unfinished Frontier

PCLR established the first U.S. organized-market architecture that places provisional large-load consumption directly inside nodal security-constrained economic dispatch while preserving a separate firm planning allocation. It separated conditional consumption from firm planning allocation. It moved incremental delivery risk toward the private party seeking earlier access and gave ERCOT a minimum sufficient reliability backstop through a bid-capping mechanism. The foundational achievement was institutional. Large-load curtailment moved from an emergency assumption into a visible, contractible, and dispatchable market obligation. That change creates room for private firming, nodal price discovery, QSE orchestration, workload controls, storage, generation, and new financing structures to compete around the same grid-facing outcome.

Why bankability remains nodal

A “naked” PCLR still carries location-specific risk. ERCOT does not provide a project-specific 8,760-hour curtailment forecast. The developer must model constraint frequency and duration, shift factors, shadow-price caps, topology changes, later interconnections, lower-voltage exposure, ramp and deviation rules, and the possibility that private mitigation will not be attributed to the load. Aurora’s local-versus-system scarcity framework discussed above, makes the diligence problem visible: curtailment can arise because power exists but cannot reach the site, because the system is short, or because both conditions occur together. A strong LPC and a network with multiple redispatch levers improve the product. A weak radial location can make it brittle. [32]

SCED compliance may itself require capital. A large load expected to move hundreds of megawatts within a small number of five-minute runs needs controls, energy storage, segmented infrastructure, or deliberately dispatchable process or workload. The price of flexibility includes the infrastructure required to make the promise real.

A serious PCLR diligence model therefore begins with the node, not a generic curtailment percentage. It maps monitored constraints and plausible future constraints, estimates helping and harming shift factors, reviews maximum shadow-price caps by voltage class, tests outages and topology changes, models the ABC persistence rule, and examines which generation resources can actually redispatch. It then overlays the site's minimum operating load, workload value, battery duration, generator availability, restoration limits, and contractual obligations.

The model should also separate frequency of curtailment from economic severity. A brief 20 MW reduction may be operationally trivial. A sustained dispatch to LPC can be existential for a facility whose irreducible operating load sits above the firm floor. The right LPC is therefore not simply the maximum allocation ERCOT eventually offers in yearly reviews as loads ramp at a selected site. It is the floor around which the customer can build a credible operating and financing plan, and that can look different, and better, than what the grid can promise or ever deliver on in firm allocations.

PCLR is most bankable where the grid-approved floor covers the site's irreducible operations and the provisional band funds incremental growth. It is less bankable where the entire commercial proposition depends on uninterrupted MPC. The framework makes the value of good siting, private firming, and workload orchestration design much easier to measure.

A STEM Policy Moat for AI Loads

The moat for AI loads, firming-ready industrial loads, and aggregated dispatchable loads alike, sits in the architecture required to make dispatchable load dependable for the power system. This is the particular value of technical policy and why innovators should spend time writing it. Once the rules are written to enable real grid-facing obligation, private technology has something concrete to solve for and finance. PCLR opens commercial space for dynamic firming, nodal load control, QSE orchestration, constraint analytics, telemetry, switching, storage, onsite generation, workload management, and performance-backed service agreements. The market rule gives those technologies a job to do and a performance standard against which customers can buy them.

ERCOT does not have to choose the battery duration, generator technology, workload profile, switching sequence, or control software that should sit behind a PCLR. The customer can combine them according to cost, reliability, duration, recovery time, and tolerance for curtailment. Five years from now, the best private firming stack may look very different from the one available today. The grid-facing obligation can remain stable while the technology beneath it gets better.

The next design work

The next stage is about turning curtailment exposure into a better-understood and more financeable risk. The architecture leaves room for solutions to improve, work that Luminary and its collaborative alliances are deeply engaged with.

For starters, the dynamic bid capping ABC mechanism should be evaluated against the full generation and storage response stack so that PCLR curtailment occurs in the economic order the market intends – and is consistent with the desired outcomes we want for scaling this solution outside of Texas and into utility areas as proposed in the Grid 2 RFC. Private mitigation needs attribution where the physics supports it. If a battery or generator paid for by one load materially relieves the constraint causing that load's curtailment, the commercial framework should preserve enough of that value for the investment to function as a hedge. [32]

Performance rules can become more precise as well. Large electronic loads often have asymmetric operating capabilities: very fast downward response, followed by slower restoration as workloads, cooling systems, batteries, generators, and switching arrangements return to normal conditions. SCED can accommodate that reality without pretending that the two directions have identical operational characteristics.

Better information may be even more valuable than another mandatory product feature. A developer considering hundreds of millions of dollars of power infrastructure should be able to understand which constraints materially affect its site, the relevant shift factors, how often ABC has been triggered, how deeply comparable Resources have been curtailed, and how those exposures change with topology and transmission additions.

Batch Zero has certainly supplied the first system-wide implementation vehicle, but the underlying product should outlive it and take the shape of permanence in market design for prosperous electricity-to-society infrastructure. Planning, interconnection, telemetry, QSE operations, and SCED can ultimately form a recurring pathway for conditional large-load access rather than transitional accommodation for one study cohort: it should not be lost on anyone that the public support record for PCLR includes industrial manufacturing loads, not just AI loads or currency miners. The internal control layer can mature alongside it. Grid 2.0 and related work show a path toward standardizing how workloads, buffers, and sources respond beneath the PCLR Base Point while ERCOT retains a single grid-facing Resource and a single dispatch obligation. [21]

The mature product should allow a project to connect with controllability, measure its exposure, buy protection against the risks worth hedging, and convert more of its load to firm service as transmission becomes available. That progression rewards customers for bringing useful operating capability to the grid today while preserving the long-run investment in shared transmission required to support tomorrow's demand.

Conclusion

This paper reviews how PCLR joined two technical histories that had developed separately between 2024 and 2026. NPRR1188 made individual Controllable Load Resources nodal. Their location, Energy Bid Curves, consumption, and effect on transmission constraints became part of SCED. Luminary's large-load and storage work in 2025 exposed the missing connection to planning: ERCOT could describe how a future controllable load would behave once registered as a Resource, yet the interconnection process had no durable way to use that enforceable operating behavior to support earlier energization. The missing storage-and-load study case made the disconnect concrete. The subsequent proposal of PGRR134 carried the concept into the Planning Guide, and PGRR145 and NPRR1325 embedded the core design in Batch Zero and the Nodal Protocols. [26,3,4,5,1]

The central insight is that this is the closet marker of success we have today to let planners have their planning certainty and let private parties take on the risk of grid headroom-enabled operational consumption at higher baseline quantities. In all U.S. markets for wires, similar solutions are needed to enable the race to AI to have real winners: sophisticated power grid customers must have a choice to consume above where the system is built today, as long as they carry the risks of those additional megawatts remaining exposed to congestion, stability limits, and real-time dispatch. The customer can choose how much of that exposure to absorb and how much to protect through storage, onsite generation, workload control, switching, contracts, or other private infrastructure. Large loads with heavily capitalized business models and massive VoCL exposure are the best customers for grid asset solutions like PCLR, and their participation supports the growing dynamics of networks as connecting more loads becomes challenging for U.S. power grids.

ERCOT had unusual advantages in getting to this point. Nodal SCED already measures the locational consequences of injections and withdrawals. Generator-side connect and manage already demonstrates that physical interconnection and full transmission deliverability need not arrive at the same moment. ERCOT's market routinely places congestion and operating risk on private Resources. PCLR carries that logic to consumption with edits that reflect the unique economic and policy position of loads in how the grid is funded and ultimately operated. However, the simplest iteration of the problem has been solved, and can be carried everywhere: a load can connect before every future transmission upgrade is complete because the megawatts above its firm floor remain under dispatch. Any system with a sufficiently detailed network model, enforceable telemetry, credible security-constrained commitment and dispatch, and an interconnection process capable of distinguishing firm from conditional service can ask the same engineering questions we did in Texas over the past year and solve with an open-access answer that invites AI loads and other loads to become grid assets. The next five years will show how quickly other markets learn these lessons.

Load growth is certainly arriving faster than transmission can be built. In all markets where this remains true, we expect system planners and operators to gravitate to the same fundamental:

"The grid provides the floor. The private sector firms the rest."

– Arushi Sharma Frank, July 14, 2026, Luminary Strategies [Link].

Citation Notes

1. Electric Reliability Council of Texas (ERCOT), PGRR145, "Batch Zero Process for Large Load Interconnections," approved June 18, 2026; ERCOT, NPRR1325, "Related to PGRR145, Batch Zero Process for Large Load Interconnections," approved June 18, 2026. PGRR145 and NPRR1325 define the adopted Batch Zero and PCLR architecture; most Batch Zero provisions became effective July 11, 2026, while specified PCLR SCED, bid-cap, ramp-rate, and energization provisions awaited system implementation. [PGRR145] [NPRR1325].

2. Arushi Sharma Frank, "Making Large Loads and Small Loads Grid-Visible and Dispatchable: How SCED Works and Why It Works," Luminary Strategies, Aug. 2, 2025 [Substack Link]. The paper explains how state estimation, contingency analysis, SCED, telemetry, reserve management, and look-ahead commitment form the operating architecture through which dispatchable load can participate in ordinary system operations.

3. Arushi Sharma Frank and Samuel Brandin, “Comments in Support of Proposed PGRR re: Controllable Load Resource Planning Ahead of NPRR1188 Implementation,” Sept. 18, 2025, pp. 1–7. [PGRR134 public record link]

4. Arushi Sharma Frank, “Controllable Load Resource Planning Ahead of NPRR1188 Implementation—Discussion Draft,” Sept. 18, 2025, pp. 1–10. This document is referred to in this paper as Version Zero. [PGRR134 public record link].

5. ERCOT, PGRR134, “Interconnection Studies Reform for Dispatchable Loads,” posted Nov. 1, 2025, sponsored and written by Luminary Strategies, later withdrawn after the Batch Zero package absorbed the underlying architecture. [PGRR134 public record link].

6. Joshua D. Rhodes and Michael E. Webber, “The Texas Flex: How Connect and Manage Can Deliver Speed to Power for Generation and Load,” GridLab and IdeaSmiths, July 2026, pp. 22–23. The paper is useful as a contemporaneous description of Batch Zero’s inclusion, system-wide study, commitment, refinement, and transmission-plan stages; its PCLR discussion largely reflects March–April 2026 materials and predates the final PGRR145/NPRR1325 architecture.

7. ERCOT, Market Notice M-B063026-01, “Partial Implementation Details of PGRR145/NPRR1325, Batch Zero Process for Large Load Interconnections,” June 30, 2026; ERCOT PGRR145 and NPRR1325 issue pages. [ERCOT Market Notice].

8. ERCOT, PGRR145, “Batch Zero Process for Large Load Interconnections,” and NPRR1325, “Related to PGRR145, Batch Zero Process for Large Load Interconnections,” approved June 18, 2026; ERCOT Large Load Integration and Form W materials. The final architecture distinguishes the PCLR’s grid-backed planning floor, bounded by the amount ERCOT determines can be reliably served for the applicable year, from the conditional operating quantity between LPC and MPC, while requiring the full requested load to remain visible in stability assessment and long-term transmission planning. [PGRR145] [NPRR1325] [Large Load Integration / Form W].

9. ERCOT, Large Load Interconnection Batch Study Workshop #7, Apr. 9, 2026, “Provisional CLR (PCLR) – Dynamic Bid Capping Process,” slide 13. The public workshop materials describe the prior-run shadow-price trigger, persistence into the current SCED run, the helping shift-factor screen of ≤ −0.02, calculation using Step 1 System Lambda, and application of the Adjusted Bid Cap before SCED Step Two. [ERCOT Workshop #7].

10. Siddharth Suryanarayanan, “Flexible AI Data-Centers Are Not Automatically Grid Dispatchable Resources,” Medium, July 24, 2026, https://moonbowtech.medium.com/flexible-ai-data-centers-are-not-automatically-grid-dispatchable-resources-c9ec333a4109. The article distinguishes technical flexibility from dispatchable Resource status, explains the boundary between commitment and interval dispatch, and identifies magnitude, ramping, duration, recovery, dependencies, telemetry, verification, and fallback behavior as features that determine deliverability. [Public article link].

11. Fred C. Schweppe, “Power Systems ‘2000’: Hierarchical Control Strategies,” IEEE Spectrum 15, no. 7 (July 1978): 42–47, especially pp. 42–46. Schweppe placed customer generation, customer storage, real-time models, spot prices, interruptible rights, communications, and hierarchical physical control inside one energy-marketplace architecture.

12. Sharma Frank, Arushi, “Making Large Loads and Small Loads Grid-Visible and Dispatchable,” *supra* note 2, at pp. 4–14.

13. Conleigh Byers and Farhad Billimoria, “Participatory Demand and New Large Loads in Electricity Markets,” Harvard Kennedy School and Oxford Institute for Energy Studies, March 2026, pp. 2–7. [Belfer Center policy brief].

14. Farhad Billimoria and Conleigh Byers, “Bits-to-Watts: Connecting Markets and Prices for Compute and Power,” working paper, July 23, 2026, especially pp. 1–7, 11–12, and market-design discussion. The paper develops Value of Compute Load, bid-in demand, full-strength price formation, non-firm open access, and hedging around efficient spot prices.

15. Farhad Billimoria and Conleigh Byers, “The Value of Compute Load: What Cloud Markets Reveal About the Potential for Flexible Data Centers,” The Power Game, July 31, 2026, p. 2. The authors expressly identify work by Arushi Sharma Frank and others as a contemporary initiative carrying forward Schweppe’s vision of demand bidding flexibility and being dispatched through SCUC and SCED. [Substack Link].

16. Conleigh Byers, LinkedIn post announcing the Value of Compute Load working paper, Aug. 4, 2026, source text archived by the author. Byers describes VoCL as the number that should sit at the core of non-firm interconnection or flexibility agreements because it determines both opportunity cost and the real-time incentive to comply with operator instructions. [Linkedin Link].

17. Rhodes and Webber, “The Texas Flex,” pp. 11–16, 22–25, and 29–32. Cited here for ERCOT’s generation-side connect-and-manage history, Batch Zero background, and the proposition that transmission expansion remains necessary alongside provisional access.

18. Luminary Strategies comments filed in the PGRR145 record on April 23 and May 18, 2026; ERCOT PGRR145 issue page, documents 145PGRR-54 and 145PGRR-108. [PGRR145 public record link]

19. ERCOT, Market Notice M-B063026-01, June 30, 2026. The notice deferred Protocol §§ 4.4.9.4, 4.4.9.4.4, portions of 6.5.7.3, and 6.5.7.11, along with Planning Guide § 9.6.1(2)–(3), to later system implementation. [ERCOT Market Notice link]

20. Frank and Brandin, Sept. 18, 2025 Explanatory Comments, pp. 2–6: nodal dispatch under NPRR1188, private firming, early energization ahead of the Load Commissioning Plan, preservation of firm planning, and privately managed delivery risk. [PGRR134 public record link]

21. J. Chris Shelton, Brett Galura, AJ Hall, Varun Joshi, and Gregory Phillips, “Introducing Grid 2.0: An Open Standard for Grid Service Access and Allocation,” Grid 2.0 Task Force, Request for Comments No. 1 (Aug. 2026), especially pp. 3–11 and 14–18. The final RFC proposes a voluntary one-minute orchestration layer between sub-second grid controls and five-minute market dispatch, using Use, Buffer, Source, and Allocator elements and the G2P, GCAP, and DESP protocols; it expressly identifies ERCOT PCLR as an implementation pathway. Supporting organizations include AES Next, LLC, Fluence Energy, Inc., and Tapestry at Google. Arushi Sharma Frank is acknowledged among the contributors; Luminary served as a contributing reviewer. [Grid 2.0 Task Force link].

22. ERCOT, Nodal Protocol Revision Request, *Related to PGRR145, Batch Zero Process for Large Load Interconnections* (Posted March 4, 2026). This Nodal Protocol Revision Request (NPRR) defines essential terminology related to the transitional Large Load interconnection process, the Batch Zero Process, and modifies the Regional Planning Group procedures to allow for a new category of projects that result from the Batch Zero Process. [Final Version, Board Report, June 2026 link].

23. Arushi Sharma Frank, “The Grid Provides the Floor. The Private Sector Must Firm the Rest: Data Center ‘Flex’ Needs New Language Fit for Purpose—Dynamic Firming Is Speed to Power,” Luminary Strategies, July 14, 2026, pp. 1–13. The article distinguishes baseline firming, operational headroom access, and attribution; defines the firm floor, dispatchable band, private firming stack, return-to-floor obligation, and attribution layer; and supplies the operating-band figure reproduced in this paper. [Luminary Strategies Substack link]

24. Agentic Infrastructure, “Dispatchable Loads,” presentation to the ERCOT Planning Load Working Group, Nov. 18, 2025, slides 2–9. The presentation maps CLR election and registration into the LLIS process, states that the load assumes all interim curtailment risk, distinguishes thermal constraints manageable through

real-time curtailment from voltage and stability upgrades that remain prerequisites, and describes PGRR134 as a technology- and design-agnostic speed-to-power pathway. [ERCOT PLWG public record]

25. Sharma Frank, Arushi. "PGRR134 (As Amended Nov. 14, 2025): Sponsor Overview," Luminary Strategies, presentation to the ERCOT Planning Load Working Group, Nov. 18, 2025, slides 1–3. The presentation assigns private POI studies and provisional delivery risk to the ILLE, distinguishes thermal violations that may be mitigated through CLR election from voltage, stability, short-circuit, and facility-study requirements, and preserves CLR status until firm Load Commissioning Plan upgrades are complete. [ERCOT PLWG public record]

26. Electric Reliability Council of Texas (ERCOT), NPRR1188, "Implement Nodal Dispatch and Energy Settlement for Controllable Load Resources," posted June 27, 2023, especially pp. 1–5 and 12–16. ERCOT sponsored the revision in response to Phase 1 of the PUCT market-design blueprint. The filing assigns individual CLRs a Resource Node, nodal shift-factor dispatch and nodal settlement, Resource-specific Energy Bid Curves, separate visibility of CLR consumption, and revised SCED treatment; ERCOT states that nodal dispatch is essential for efficient congestion management and crucial for reliable operation as large-load size increases. [NPRR1188 public record]

27. Lancium LLC (comments filed by Eric Goff and Andrew Reimers), Comments on NPRR1188, Oct. 5, 2023, especially pp. 1–3. Lancium proposed using net Load from the grid rather than total CLR Load for Load Ratio Share settlement, argued against economic disincentives to CLR registration for co-located Load, and identified interactions with NPRR1191 and 4-Coincident Peak treatment. [NPRR1188 public record link].

28. ERCOT (Dave Maggio), Comments on NPRR1188, Apr. 4, 2024, especially pp. 1–2. ERCOT agreed that co-located CLRs should receive the same treatment as other co-located Load, separated that issue from the broader question of transmission-cost allocation, and restored the ability of a CLR to update its Energy Bid Curve at any time. [NPRR1188 public record link].

29. Oncor Electric Delivery Company LLC (Martha Henson), Comments on NPRR1188, July 15, 2024, especially pp. 1–2. Oncor focused on the practical engineering, access, billing, and consent issues associated with behind-the-meter CLR settlement metering while preserving the utility's obligation to serve. [NPRR1188 public record link].

30. ERCOT, NPRR1188 Board Report, Oct. 10, 2024, especially pp. 4–7; ERCOT, NPRR1188 PUCT Report, Nov. 21, 2024, especially pp. 5–7. The record notes support at PRS for expanding nodal treatment to all large Loads rather than only CLRs; unanimous recommendations by PRS, TAC, and the ERCOT Board; Independent Market Monitor support; and PUCT approval on Nov. 21, 2024. [NPRR1188 public record].

31. Agentic Infrastructure, LLC and Luminary Strategies, LLC, Joint Commenters' Comments to PGRR134, Nov. 14, 2025. In revising PGRR134, the sponsors expressly separated thermal constraints manageable through CLR redispatch from IROLs and other non-thermal limits—including voltage and stability criteria—that could remain conditions precedent to provisional energization and require completion of identified ERCOT/TSP projects. [PGRR134 public record] [Joint Commenters comments].

32. Aurora Energy Research, "SCED Model Assessment: How Local Transmission Constraints and System-Wide Scarcity Can Drive PCLR Curtailment," figure provided to Luminary Strategies Aug. 10, 2026, from expected Aurora Energy Research publication Aug. 26, 2026. The assessment separates PCLR curtailment risk into local transmission congestion and system-wide scarcity and illustrates the four combinations of those conditions.

33. Sharma Frank, Arushi. "Dispatchability, Not Vibes: What PJM Market Monitor Gets Right—and Misses—on Data Centers Flex." Luminary Strategies, Nov. 17, 2025. [Substack article link].

34. Sharma Frank, Arushi. “The Texas Energy Reference Design: Stress-Testing Load Growth Challenges for an AI Century.” Center for Strategic and International Studies, Apr. 7, 2026. [CSIS commentary].

35. Public Utility Commission of Texas, Project No. 58317, SB 6 Implementation Workshop, July 21, 2025, remarks of Arushi Sharma Frank, including discussion at approximately 3:12:49 and 3:37:11; see also Arushi Sharma Frank, “From Bottlenecks to Breakthroughs: Rewriting the Grid Planning Playbook in the Southwest Power Pool and Texas,” Center for Strategic and International Studies, July 23, 2025. The contemporaneous analysis describes Texas’s missing planning-phase pathway for non-firm large-load transmission access and the need to make controllable-load and paired-resource behavior usable in transmission studies where the operating system can enforce it. [PUCT SB 6 workshop record] [CSIS contemporaneous analysis].

36. Energy Central Podcast, Feat. Arushi Sharma Frank, interview by Kinsey Grant Baker, “How to turn AI into a power plant with Arushi Sharma Frank,” Power Perspectives, Energy Central, Nov. 12, 2025, video podcast, especially discussion beginning at 13:23. The discussion addresses the distinction between flexibility and dispatchability, grid-visible orchestration of large-load resources, and the need to translate behind-the-meter or campus-level capability into a firm, measurable grid-facing response. [Exact video timestamp] [Energy Central episode].

37. Sharma Frank, Arushi, in interview by Oliver Kerr, “Can Data Centres Earn Their Power?,” Energy Unplugged by Aurora, Episode 299, Aurora Energy Research, July 28, 2026. [Aurora Energy Research episode]

38. Myers, Liz. “Infinium Energy: Industrial Ultra-Low Carbon Fuels.” Presentation to ERCOT Large Load Interconnection Batch Study Workshop #3, Feb. 26, 2026. [ERCOT presentation PDF link].

39. PGRR134 drew support from a broad cross-section of power-market, infrastructure, technology, development, and large-load stakeholders. An ad hoc coalition supporting Luminary Strategies’ proposal for dispatchable-load architecture included signatories Aaron Smith, Point Three LLC; Andrew Schaper, Schaper Energy Consulting, LLC; Aroon Vijaykar, Emerald AI; Arushi Sharma Frank, Luminary Strategies, LLC; Brian Kauffman, Mainspring Energy, Inc.; Clayton Greer, Cholla Petroleum, Inc.; Doug Lewin, Stoic Energy, LLC; Ed Scarborough, Distributed Sun LLC; Jeff Bladen, Verrus; Jigar Shah, in his personal capacity; Jonathan Abebe, Cloverleaf Infrastructure; Kevin Boudreaux, Monarch Energy; Liz Myers, Infinium Operations LLC; Luke Hansen, Clean AI Energy; Lynne Kiesling, Knowledge Problem LLC; Matthew Crosby, Cypress Creek Renewables; Peter Hans Hirschboeck, impactECI; Phil Harris, former Chairman, President and CEO of PJM and PharrisLLC; Sam Brandin, Agentic Infrastructure, LLC; Sean Voigt, Energy Innovation Hub TX; Timothy Hade, Brightfield Infrastructure, LLC; Tim Profeta, Novi Strategies LLC and co-author of Rethinking Load Growth: Assessing the Potential for Integration of Large Flexible Loads in US Power Systems, published by Duke University’s Nicholas Institute for Energy, Environment & Sustainability in 2025; and Val Angelkov, Engelhart Commodities Trading Partners (US) LLC. The formal PGRR134 record also accumulated separate comments and technical engagement from Cypress Creek Renewables, impactECI, GridCARE, Camus Energy, Oncor Electric Delivery Company, and V-Wire, in addition to earlier comments from Agentic Infrastructure, Schaper Energy Consulting, and ERCOT. See ERCOT, PGRR134, Interconnection Studies Reform for Dispatchable Loads, Key Documents 134PGRR-02 through 134PGRR-17, Nov. 1, 2025–Feb. 6, 2026. The Reliability and Operations Subcommittee subsequently tabled the proposal and referred the issue to the Planning Working Group for further development; PGRR134 remained under active PLWG consideration into spring 2026 before its May 27 withdrawal following ERCOT’s development of the Batch Zero package which incorporated the market design proposal. [PGRR134 public record] [Duke study].

40. Special thanks to collaboration and ideas-exchange partners including Analog Devices, Baker Botts, Brattle Group, Cholla, Inc., CIGRE, Electric Power Engineers, Enchanted Rock, ESIG, Generate Capital, Goff Policy, the GRID 2.0 Task Force (Chris Shelton et al.), Institute for Humane Studies, Mintz Levin, NERC, NRG, United Electric Cooperative, Nvidia, Oncor Electric Delivery Services, Tesla, and the Texas Competitive

Power Association. Special commendations to Pablo Vegas, Bill Flores, Woody Rickerson, Dan Woodfin, Sai Moorty, Kenneth Ragsdale, Kristi Hobbs, Matt Mereness, Jeff Billo, and many other ERCOT leaders, managers, and subject matter experts.

41. Luminary Strategies' principal and founder is an angel investor and active advisor to multiple ventures seeking to build a dispatchable, grid-forward future for large loads and distributed loads. Three of the firm's advised ventures actively participated in the PGRR145 – Batch Zero – PGRR134 proceedings cadence: Agentic Infrastructure, Emerald AI, and V-Wire. Emerald AI presented on dispatchable computational workload flexibility to LLWG on March 28, 2025 [presentation link]. Agentic Infrastructure presented on energy storage dispatch on September 19, 2025 [presentation link]. V-Wire submitted comments to LLWG on March 12, 2026 [presentation link].