GCA

Global Capacity Alliance

TELEMETRY OFFSETTING TECHNICAL ARCHITECTURE – ERCOT WHITE PAPER

Prepared By: V-Wire for the GCA Board

Date: September 2026

Status: Draft Technical Architecture

Subject: Establishing the ERCOT-embodiment of Telemetry Offsetting to enable curtailable load integration through equivalently offsetting resources.

Stakeholder Notice: The technical architecture and proposed protocol concepts outlined in this white paper reflect the current operational recommendations of the Global Capacity Alliance (GCA) as of September 2026. This document is intended to serve as a baseline for discussion within the ERCOT stakeholder process.

Executive Summary

The Brattle Group recently identified the preferred pathway to fix the large load interconnection bottleneck[1]: netting flexible loads with co-located generation to bound net grid withdrawals. **Telemetry Offsetting** is the ERCOT embodiment of this nationally recommended solution, providing a software-defined framework for netting a flexible load's curtailment obligation with a paired co-located resource, whether in a Front-of-the-Meter (FTM), Behind-the-Meter (BTM), or other shift-factor neutral configuration. The logical progression below proves the operational and commercial necessity of this framework:

1. **Speed requires Flexibility.** Connecting massive compute load years ahead of transmission upgrades requires flexible connections.
2. **Flexibility requires Generation.** To be bankable, large loads must privately mitigate curtailment risk through new Bring-your-own-Generation (BYOG).
3. **Generation requires Attribution.** Capital will only fund that new generation if the market engine actually credits it to the load, providing a proportional reduction in curtailment exposure.
4. **Market Rules Force Separation.** Regardless of how a site is physically wired—whether FTM across dual POIs or BTM within a PUN—grid operator rules rightfully mandate separate Resource IDs and distinct telemetry. Dispatched independently, the grid engine (SCED) defaults to a gross obligation that socializes the generator's relief to the grid while simultaneously curtailing the load, creating the "attribution gap".
5. **Telemetry Offsetting is the Tether.** Telemetry Offsetting is the software implementation layer. By dynamically tethering operational limits via existing telemetry managed by the asset's Qualified Scheduling Entity (QSE), SCED sees the paired physical boundary, securing a net curtailment obligation while providing unblended operator visibility, integrating new generation, and protecting the flexible load.

New load needs new generation. Driven by the load's need for speed, that generation is coming. Bridging the attributability gap ensures gigawatts of privately funded generation are reliably integrated into the ERCOT market rather than pushed off-grid. This paper provides the technical details of the solution: Telemetry Offsetting.

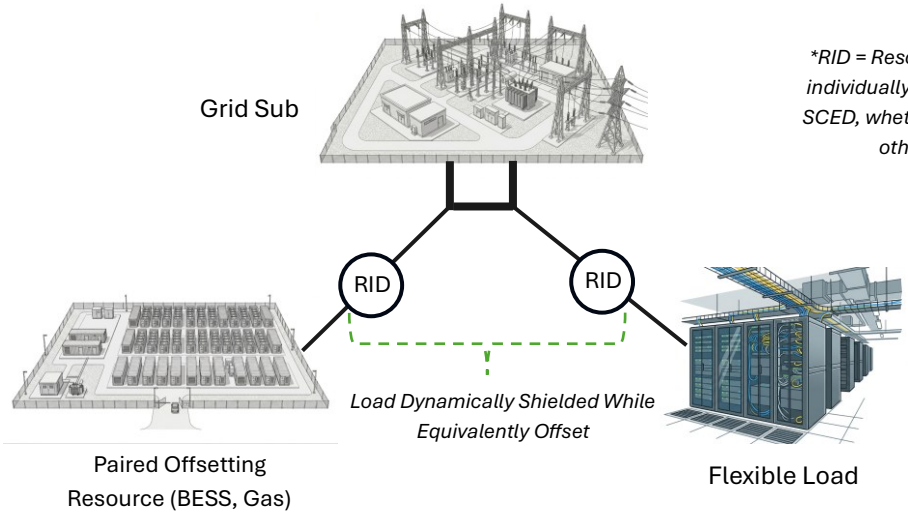

**RID = Resource ID, equating to an individually dispatched resource in SCED, whether in FTM, BTM, PUN or other configuration*

[1] Brattle Report: Solutions for Timely Interconnection of Large Loads 15 Sept. 2026 Prepared for Google

A Note on Global Applicability & Topological Neutrality

While this white paper utilizes ERCOT as its foundational case study, the attributability gap is a universal flaw in nodal electricity markets. The Telemetry Offsetting architecture proposed herein is fundamentally market-agnostic. Furthermore, it is a software-defined, configuration-agnostic architecture. Whether a facility utilizes Front-of-the-Meter (FTM), Behind-the-Meter (BTM), or Private Use Network (PUN) configurations, the mechanism functions identically. By leveraging distinct Resource IDs and unblended telemetry streams, Telemetry Offsetting resolves localized transmission constraints while accommodating any shift-factor neutral physical topology.

1. The Basics: PCLRs and the Firm Floor

When a data center interconnects as a Provisional Controllable Load Resource (PCLR)[2], ERCOT's Security-Constrained Economic Dispatch (SCED) respects its telemetered **Low Power Consumption (LPC)** limit as a hard curtailment floor. Any operational MWs above this floor are considered by SCED to be curtailable.[3]

Private Mitigation Options: The Attributability Gap and the Telemetry Offsetting Solution

A large load operating as a PCLR may try to pair a new resource (such as BESS or thermal) to help manage expected load curtailment. Regardless of whether the load and resource sit behind a shared POI or within a single co-located facility boundary but with two POIs, SCED cannot natively link them. Because SCED treats each as an independent asset, the resource's mitigation is effectively socialized to the grid while the load remains exposed to curtailment. This risk has been studied by Aurora Energy Research and further documented by Brattle in recent public reports.

Telemetry Offsetting bridges this "attributability gap" using existing market concepts and leverages existing operational pathways. Targeted protocol changes have been drafted for stakeholder consideration that will allow a PCLR to mitigate curtailment risk within normal SCED operations through 3 core principles:

1. **The Telemetry Shield:** The Qualified Scheduling Entity (QSE) dynamically raises the PCLR's LPC telemetry setpoint to match the paired resource's real-time market availability. This shields the load from SCED curtailment—but only while the offsetting capacity is **actively offered and equivalently offsetting**. If the resource exhausts or trips, the shield drops.

2. **ERCOT Verifies:** Post-operating day settlement occurs natively using the telemetry and meter data ERCOT already collects for both resources, proving that every shielded MW was backed by the paired resource, the same way it already validates Aggregate Load Resource telemetry today.

3. **Deterministic Macro-Scarcity Yielding:** While protected during localized congestion, the architecture enforces mandatory yielding during system-wide emergencies. If the grid enters macro-scarcity, the QSE drops the shield, releasing both the resource and the load to SCED.

Through this framework, Telemetry Offsetting provides a path for large loads to utilize the existing PCLR framework while bringing their own new capacity to privately mitigate their risk while upgrades are built all within the existing transmission planning framework. Once the required transmission upgrades are built, the Telemetry Offsetting pairing contract ends, the load is firm and the resource continues as a full market participant, already operationally connected on an independent, grid-connected position.

---

[2] See Arushi Sharma Frank, "Flexible Loads in ERCOT = Provisional CLRs," Luminary Strategies (April 9, 2026), outlining how ERCOT's interconnection process allows large loads to split their requested megawatt capacity into a transmission-planned firm slice (LPC) and a SCED-dispatchable flexible slice

[3] See Arushi Sharma Frank, "The Power Grid 'Headroom' Paradox and Real Limits on Batteries," Teach What I Know (July 26, 2026), detailing ERCOT's structural separation of conservative Low Power Consumption (LPC) firm floors from Maximum Power Consumption (MPC) dynamic ramps to safely unlock grid headroom for flexible assets.

The Attributability Gap: Example

As shown in Figure 1, consider a 200 MW large load operating 100 MW firm and 100 MW as a PCLR co-located with a 100 MW BESS built to mitigate the curtailment risk to the PCLR component of the load.

1. **Without Telemetry Offsetting:** Because ERCOT assigns them separate Resource IDs, SCED views them as fully independent. During localized congestion, SCED will dispatch the 100 MW battery to help the grid while simultaneously curtailing the flexible 100 MW portion of the load. The developer financed the BESS to protect their data center, but the grid "socialized" the battery's power leaving the data center still subject to SCED curtailment. Independent modeling by Aurora Energy Research proves this architecture works, demonstrating that netting effectively closes the attributability gap and reduces forced curtailment hours by 83%.[4]

2. **The Solution (With Telemetry Offsetting):** The QSE managing both assets utilizes real-time telemetry to create a Telemetry Shield. Because the resource is bid into the market, the QSE raises the data center's Low Power Consumption (LPC) limit on an equivalent basis by 100 MW (to 200 MW given the example), effectively shielding the otherwise curtailable portion of the load in SCED's eyes.

3. **The Result:** SCED sees the data center's 200 MW LPC as a hard floor and bypasses it for curtailment. Simultaneously, SCED dispatches the 100 MW battery to resolve the congestion. The battery provides physical, **equivalent** counter-flow, the grid remains safe, and the large load operates without interruption. **The BESS mitigation is now directly attributed to the load.**

Figure 1: Telemetry Offsetting Diagram

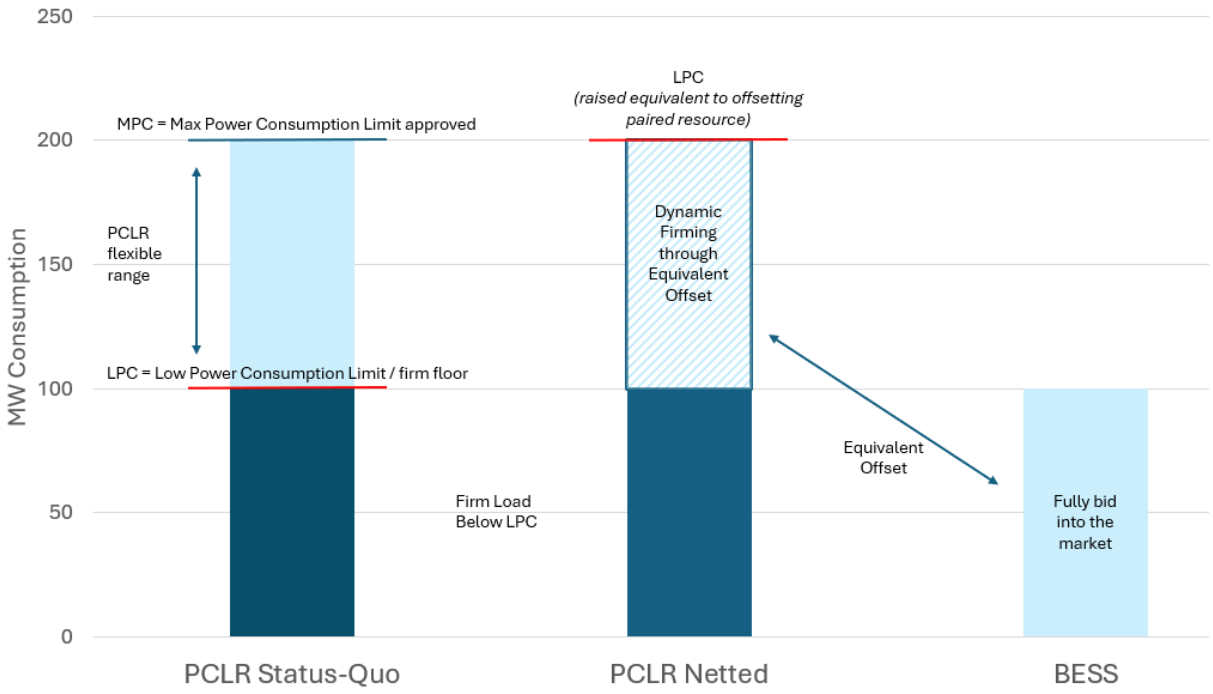

[4] Aurora Report press release and Report: Curtailment Mitigation for PCLRs: Modeling the Effects of Front-of-Meter Netting 15 Sept. 2026 prepared for Google and V-Wire. Note an FTM configuration was assessed but the results hold for any configuration using co-located, separately dispatched assets (BTM, PUN, FTM, etc.).

2. Operations: Paired Resource Bid Capping

To successfully function as proxy for the paired PCLR, the paired generation asset must reliably clear the merit order on an economically equivalent basis as the PCLR. To achieve this, the paired resource is subjected to the same Automatic Bid Capping (ABC) methodology currently grey-boxed for PCLRs.

Under this architecture, ESRs are mitigated to the nodal ABC, while Generation Resources (thermal) are mitigated to the lesser of their cost-based Mitigated Offer Cap (MOC) or the nodal ABC (min(MOC, ABC)). Because the generator's cap will always be equal to or lower than the load's curtailment threshold, SCED will dispatch the resource to relieve nodal congestion equivalently to the PCLR had it not been shielded.

3. Operations: Transmission Planning Alignment & The Scarcity Release

A critical advantage of the Telemetry Offsetting architecture is that it does not require structural changes to ERCOT's Regional Planning Group (RPG) methodologies or long-term base-case models. Telemetry Offsetting is integrated into Real-Time Operations through existing telemetry, applied to two independently dispatched and modelled resources. Because it does not alter the physical topology of the grid, it maintains alignment with ERCOT's existing planning assumptions:

- **Distinct Interconnection Studies and Resource Modeling:** From a transmission planning perspective, ERCOT rules mandate that the PCLR and the mitigating generation resource (BESS, Thermal, etc.) are studied and modeled independently, regardless of their physical POI configuration. Under the rules established by NPRR1188, a CLR must be metered and telemetered separately from co-located generation. Consequently, the mitigating generation must proceed through the standard Generator Interconnection or Modification (GINR) process, while the load proceeds through the Large Load Interconnection Study (LLIS) process (e.g., Batch Zero). This ensures that ERCOT planners establish independent baseline parameters, stability models, and Resource IDs for both assets before they are eligible for Telemetry Offsetting in real-time operations.

- **Local Congestion & Operational Equivalency:** Long-term transmission planning models rely on nodal injections and withdrawals to evaluate physical line limits under both N-0 steady-state and N-1 contingency conditions. The core operational principle of Telemetry Offsetting is maintaining a net-zero physical impact on a constrained node. While both the PCLR and the mitigating resource remain fully modeled, independently telemetered, and visible to ERCOT in all configurations, Telemetry Offsetting establishes the mechanism to privately sequence their market response to achieve this net-zero state at the node. It ensures that the privately financed resource is dispatched as the first line of defense to provide the required counter-flow, and the PCLR curtails if the paired resource exhausts its duration. Because the physics of load-flow modeling are governed by the net power flow at the Resource Node, the offset operational sequencing ensures line limits remain securely bounded.

- **The Scarcity Release:** During system-wide coincident peak or scarcity scenarios, the Telemetry Offsetting protocol executes a "Scarcity Release." The QSE drops the PCLR's LPC limit to the firm floor, unshielding the flexible portion of the load. Consequently, both assets behave as ERCOT Planning models standard independent PCLRs and resources during grid emergencies.

Because of this alignment, Telemetry Offsetting allows developers to execute commercial attributable links in operations while allowing ERCOT Planners to continue utilizing their existing study methodologies.

4. Operations: Execution of the Scarcity Yield and Emergency Operations

To execute the Scarcity Yield, the architecture relies on a staggered, deterministic sequence of market-driven obligations to ensure capacity is available during grid scarcity events.

Triggering the Yield (Deterministic Scarcity & Operational Unshielding)

Rather than relying on proprietary QSE telemetry estimations or bespoke shadow-price logic, the Telemetry Offsetting architecture executes a scarcity yield through two strict protocol layers.

- **The Scarcity Yield (EEA 1+):** The telemetry shield is suspended during system-wide emergencies. The Scarcity Yield is not triggered by a QSE's internal forecast, but by ERCOT's issuance of an Extensible Markup Language (XML) message or ERCOT Hotline call declaring an Energy Emergency Alert (EEA) Level 1+. Upon declaration, the QSE must, within 60 seconds and without intentional delay, drop the telemetered LPC to its unshielded Baseline LPC. ERCOT measures the elapsed time between its own timestamped declaration and its receipt of the telemetry drop.

- **The Physical Capability Yield (Outages and Exhaustion):** If the paired mitigating resource loses the capability to inject offsetting megawatts—whether due to a Forced Outage, physical derate, or an energy storage resource exhausting its State of Charge—the QSE must reduce the telemetry shield commensurately immediately. This guarantees that no load remains shielded in SCED unless the equivalent offsetting physical capacity is live and available.

Once the load is released for potential curtailment during an emergency, ERCOT's SCED algorithm natively optimizes by need.

5. Operations: Day-Ahead Alignment and DRUC Management

To ensure the Day-Ahead Reliability Unit Commitment (DRUC) process accurately reflects the net-neutral physical reality of the Telemetry Offset pair, the framework bifurcates Current Operating Plan (COP) reporting requirements of the paired assets based on the mitigating resource's technology. This prevents ERCOT's planning engines from double-counting the pair for the coming day given their netting relationship. The architecture manages this Day-Ahead visibility through two operational pathways:

- **The Energy Storage Resource (ESR) Pathway:** When the Provisional Controllable Load Resource (PCLR) is paired with an ESR, the reporting burden falls on the ESR to prevent its reserved capacity from showing as available to the broader market. The ESR's COP reflects a High Sustained Limit (HSL) and an Hour-Beginning Planned Target State of Charge (SOC) that exclude the Available Paired Telemetry Offset Capacity. Consequently, if the entirety of the ESR's capacity is reserved to shield the co-located load, the Qualified Scheduling Entity (QSE) submits a Resource Status of ON, an HSL of 0 MW, and a Target SOC of 0 MWh. This ensures DRUC treats the reserved capacity as unavailable for system-wide dispatch while the PCLR load's capacity flags as available. This applies only to day-ahead COP submission; real-time SCED telemetry for the ESR is unaffected and continues to reflect dispatchable capability.

- **The Generation Resource Pathway:** Conversely, when the PCLR is paired with a standard Generation Resource (e.g., a natural gas facility), the mechanism inverts: the load is reported as unavailable for curtailment while the generator remains visible but commercially restricted. The PCLR's COP reflects a Low Power Consumption (LPC) limit that includes the expected offset. This effectively removes the shielded portion of the load from ERCOT's DRUC assessment of available curtailment capability. Concurrently, the paired Generation Resource submits an On-Line Resource Status with an HSL sufficient to support the offset. The specific portion of the generator's capacity committed to the Telemetry

Offsetting arrangement is ring-fenced, reported separately, and rendered ineligible for standard RUC commitments or Ancillary Service awards.

By tailoring the COP reporting to the asset class, the Telemetry Offsetting framework provides ERCOT operators with a transparent, accurate view of the node without distorting Day-Ahead price formation or system reserves.[5]

6. Operations: Navigating ESR Exhaustion & Ramp Compliance

Because the Telemetry Offsetting architecture may rely on a finite ESR (e.g. a co-located BESS) to shield the data center during prolonged nodal congestion, the facility must proactively manage the eventual exhaustion of that ESR. How the facility manages this transition can be managed in two ways:

Pathway A - Native SCED Compliance (Direct Exposure)

For highly elastic loads capable of rapid, single-interval curtailment, a complex transition is unnecessary. As the mitigating resource reaches exhaustion, the QSE simply moves the telemetered shield to the firm floor. The flexible portion of the load is exposed to SCED. Because the facility hardware can safely respond to rapid directives, it accepts and complies with standard, unmitigated SCED Base Point dispatch instructions.

Pathway B - The Coordinated Zero-Out (Synthetic Ramping for Inflexible Loads):

Every dispatched Resource — including an unmasked PCLR — remains bound by ERCOT's standard SCED Base Point-following and ramp rate requirements. For large loads that can't meet those requirements directly, the NPRR's Shared Performance Testing provision (Section 8.1.1.2(15)) enables an alternative: the PCLR may qualify against those requirements using the net telemetered response of the load and its paired resource together, rather than the load's physical response alone. Having qualified this way, the facility is expected to deliver that combined response in real time; a QSE that lets the BESS deplete passively and hands the load to SCED at the last minute is exposing the PCLR to standard Set Point Deviation Charges for failing to curtail as tested. The commercially rational alternative enabled through this pathway is for the QSE to proactively synthesize a coordinated ramp-down using the remaining battery duration:

- **The Trigger (The Telemetry Runway):** The BESS reaches a minimum State of Charge (SoC) threshold (e.g., 25% or ~45 minutes of duration remaining) while localized congestion is binding.
- **The Transition:** The QSE notifies the data center to begin a controlled, closed-transition to onsite standby generation or through shifting IT load.
- **The Symmetrical Glide Path:** Over the final 15 minutes of ESR SoC (or whatever the load requires), server load is systematically migrated off the ERCOT grid. As the facility's physical consumption decreases, the QSE proportionally lowers both the BESS injection output and the PCLR's LPC telemetry limit.
- **The Result:** The BESS injection and the grid load net out in parallel. Because the reduction in load offsets the reduction in BESS discharge, the net megawatt impact to the constrained transmission line is offset.

To formalize this coordinated dispatch and overcome the physical inflexibility of certain large-scale computing operations, the framework permits the facility to utilize newly defined Shared Performance Testing Capabilities. A PCLR is authorized to use the high-speed physical ramping and inverter control capabilities of its paired mitigating resource to fulfill ERCOT's SCED Base Point-following and ramp rate performance testing requirements. ERCOT evaluates the success of these qualification tests based on the net telemetered response at the shared Resource

[5] We recognize that certain updates are needed for PCLRs to be accurately included in the DRUC process regardless of Telemetry Offsetting's adoption.

Node. This provides a vital commercial bridge, allowing facilities with otherwise rigid hardware to successfully qualify for and operate within ERCOT's dynamic dispatch parameters.

This coordination, including SoC trigger thresholds, notification timing, taper duration, etc., is privately managed. The NPRR provides the qualifying test-coordination mechanism that makes net-response compliance possible in the first place, enabling a framework that makes private execution the rational choice over passive exposure while leaving underlying reliability and operational requirements whole.

Unplanned Outages: Compliance Protection & The Controlled Ramp Grace Period

While planned ESR depletion is managed through a coordinated glide path, an unexpected Forced Outage or physical derate of the paired resource necessitates a rapid, structured telemetry transition. Under current PCLR rules, an unmitigated Base Point deviation following a generator trip will create regulatory uncertainty regarding PCLR registration status.

The framework manages this friction by establishing a compliance safe harbor during unplanned outages:

- **Automated Telemetry Adjustment:** Upon a paired resource trip, the QSE immediately drops the telemetered LPC shield. The framework evaluates if the load was being actively shielded from curtailment—verified by whether the Real-Time Locational Marginal Price (RTLMP) equaled or exceeded the maximum price on the PCLR's Energy Bid curve during the SCED interval preceding the event.
- **EDRR-Executed Grace Period:** If the economic curtailment condition is met, a SCED instruction to curtail on the next interval is imminent. To safely manage this, the PCLR begins physically reducing its consumption within sixty (60) seconds of the event, targeting its newly reduced telemetered LPC. The reduction must be executed at its telemetered Emergency Ramp Rate Down (EDRR). To facilitate this, the facility is granted a dynamic compliance grace period. The duration is calculated by dividing the required megawatt reduction by the PCLR's telemetered EDRR, rounded up to the next whole 5-minute SCED interval. *(For example: A 100 MW required reduction at a 15 MW/minute EDRR takes 6.6 minutes; rounded up, the facility is granted a 10-minute, two-SCED-interval grace period)*.
- **Status Protection:** Base Point deviations incurred strictly within this EDRR-defined ramp window are exempt from compliance evaluation and cannot be used as grounds for status suspension or revocation, provided the final Base Point is achieved by the grace period's conclusion.

This mechanism replaces an unquantifiable compliance risk with a deterministic, physically tested ramp-down, protecting project capital while maintaining predictable system physics for ERCOT.

7. Post-Operating Day: Compliance & Settlement

Because the PCLR and its paired Resource are both directly telemetered to ERCOT, and because ERCOT already collects and retains this data, ERCOT can perform the Telemetry Offsetting compliance check natively:

1. **Volumetric Equivalency**: comparing the PCLR's telemetered LPC against the paired Resource's telemetered HSL and output to ensure the Available Telemetry Offset Capacity is equal to or greater than the Offset LPC MW for every individual SCED interval within the Settlement Interval.
2. **Latency Compliance**: Comparing ERCOT's own EEA-declaration timestamp against ERCOT's own receipt timestamp for the ICCP telemetry update reflecting the Scarcity Yield.
3. **Offer Sufficiency:** Confirming the paired Resource's Energy Offer Curve or Energy Bid/Offer Curve met the required volume.

4. **Ancillary Service Independence**: Confirming the paired Resource's telemetered High Sustained Limit (HSL) minus its Ancillary Service Resource Responsibility was sufficient to cover the Offset LPC MW, preventing the double-dipping of committed capacity.

5. **Nodal Price Alignment:** Verifying that if RTLMPs between the PCLR and paired resource diverge during any interval, the QSE adjusted the PCLR's telemetered LPC back to its Baseline LPC no later than the subsequent SCED interval.

This mirrors the Telemetry Validation methodology ERCOT already applies to Aggregate Load Resources (ALRs) under Section 22, Attachment O, the same precedent for ERCOT-run validation. Because the paired resource's SCED offer is capped at or below the load's Adjusted Bid Cap at the point of dispatch (see Section 3, Paired Resource Bid Capping), the resource is structurally incapable of clearing above the load's ABC. If ERCOT's own data indicates a QSE shielded the load without a sufficient offset, ERCOT assesses financial penalties and any other compliance penalties as necessary.

8. Planning: Structural Interconnection Architecture

A primary advantage of the Telemetry Offsetting architecture is its physical flexibility. Telemetry Offsetting is a software-defined telemetry mechanism that operates independently of physical switchyard geometry. To ensure seamless integration with ERCOT's Regional Planning Group (RPG) models, any physical implementation must satisfy one foundational engineering requirement: **Topological Equivalence at the Resource Node**.

Under Appendix A of the ERCOT Nodal Protocols, assets that connect to the same Electrical Bus map to an identical Resource Node, ensuring their shift-factor vectors across every transmission constraint in SCED are mathematically equal. Market participants can achieve this topological equivalence through two primary physical archetypes:

1. **Dual-POI Bus-Bar Branching (FTM Topology):** The load and generator execute separate drop-downs off a shared TSP switchyard bus bar. This configuration utilizes dedicated revenue meters to preserve 4CP tariff allocations and Wholesale Storage Load (WSL) tax exemptions while eliminating private-wire retail wheeling friction.

2. **Single-POI Sub-Metered Interconnection (PUN / BTM Topology):** The load and generator interconnect behind a single utility POI, utilizing internal EPS sub-meters mandated by NPRR1188 to establish distinct Resource IDs. This archetype minimizes switchyard CAPEX while remaining inherently immune to external grid sectionalization risks.

Regardless of the chosen physical archetype, both assets complete their respective study processes, ensuring ERCOT planners maintain unblended visibility into the physical attributes of each asset.

Automated Protocol Fallback (Topological Safety Net)

Regardless of the chosen physical interconnection, the Telemetry Offsetting framework includes an automated software safeguard to ensure macro-grid reliability is never compromised by local topological anomalies. Should any planned outage or unplanned physical event cause the paired assets to become electrically separated during active grid congestion, their respective Real-Time Locational Marginal Prices (RTLMPs) will diverge. Under the Telemetry Offsetting protocol, any RTLMP divergence between the paired Resource IDs serves as an automated trigger for the QSE to immediately drop the telemetry shield, safely returning both assets to standard, independent SCED dispatch.

9. The Commercial Vehicle: Telemetry Offsetting Agreement

While Telemetry Offsetting provides the regulatory and operational telemetry to keep co-located assets on the macro-grid, the architecture is ultimately powered by a private commercial instrument: the **Telemetry Offsetting Agreement**. Where possible, the Global Capacity Alliance (GCA) intends to standardize the contractual structure underpinning the agreement through an existing commercial standards body—following the precedent NAESB has set with its Base Contract for natural gas—rather than requiring a bespoke agreement framework each facility.

GCA participants and other private stakeholders are currently sizing and designing these contracts in partnership with the financial and insurance market. There is significant interest and active engagement from some of the largest insurance and financial industry actors in the world in participating in this new market construct.

Figure 2 provides an outline of this contractual arrangement of an FTM configuration. Other commercial configurations are also possible, for instance with one party playing multiple roles (e.g. resource owner + QSE).

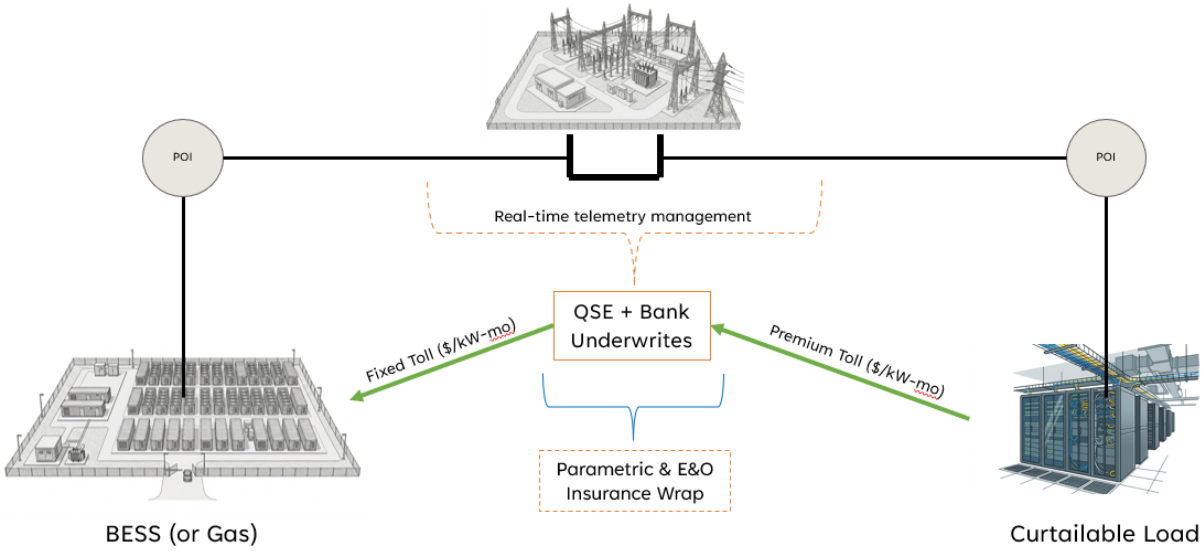

Figure 2: Telemetry Offsetting Agreement Diagram of a FTM configuration

Because ERCOT's SCED engine cannot natively link independent assets, the Telemetry Offsetting Agreement serves as the bilateral financial bridge. By keeping the commercial arrangement entirely private, the agreement successfully insulates ERCOT from dispute resolution while aligning the financial incentives of all counterparties. The agreement operates on three primary commercial mechanics:

- **The Agreement Provider (QSE + Bank):** The agreement relies on a centralized balance sheet (the "Provider") to bridge the operational needs of the large load with the financial needs of the resource. The Provider is responsible for tolling the resource's capacity and providing the dynamic firming to the load contractually and operationally through real-time telemetry management.
- **The Actuarial Wrap and Business Interruption:** Without the agreement, a PCLR faces un-warned grid curtailments, representing an uninsurable operational profile. The agreement's telemetry shield protects the load from curtailment and, when State of Charge is running low (in the case of BESS), can provide a defined X-minute (60, 120, etc.) warning before load curtailment is required. This transforms unpredictable curtailment into deterministic, planned operational expenses. The deterministic nature of

the exposure is designed to support parametric insurance structures, including congestion policies and technology performance wraps, which are currently under development with institutional underwriters.

- **The Temporal Bridge and Sunset Clause**: The agreement is a structured bridge from early energization to firm load. It allows a large load to interconnect and operate at a high commercial capacity 3 to 5 years ahead of transmission upgrades. The agreement naturally includes a "Sunset Clause". Once the Transmission Service Provider completes the necessary structural upgrades, the data center is upgraded to firm load status, the telemetry shield is not needed, and the fully amortized, co-located resource is contractually released into the merchant market.

10. Statutory Alignment: Texas SB 6 & PUCT Directives

Any structural modification to the dispatch and settlement of Large Flexible Loads in ERCOT should be evaluated against the statutory framework established by Texas Senate Bill 6. SB 6's large-load provisions are comprehensive: PURA §§39.169[6] and 37.0561[7] impose additional ERCOT/PUCT study, approval, and operational-authority requirements specifically on Behind-the-Meter (BTM) and Private Use Network (PUN) arrangements, particularly those pairing new loads with existing (pre-September 2025) generation. Telemetry Offsetting is focused on new generation, using a 12-Month Rule to ensure the paired resource is newly built rather than existing. More importantly, by mandating distinct Resource IDs and unblended telemetry for co-located assets, Telemetry Offsetting allows BTM and PUN configurations to fully comply with the broader reliability policy underlying SB 6—ensuring large loads remain curtailable and grid-visible through ERCOT's existing PCLR and SCED framework.

The Telemetry Offsetting architecture provides a software-defined framework that aligns with SB 6's goals at scale while ensuring the commercial viability of the operations.

- **Advancing the Curtailability Policy Goal (The Net-Zero Equivalent):** SB 6 dictates that non-firm large loads must be fully curtailable to protect system frequency and capacity reserves. Telemetry Offsetting is intended to advance the reliability policy goal reflected in SB 6's large-load provisions by ensuring a "Net-Zero" physical impact on the constrained node. During localized congestion, the PCLR technically remains energized, but the paired mitigating resource (BESS/Gas) injects 1:1 counter-flow. To the physics of the grid—and to SCED—the load's impact has been fully mitigated. Furthermore, the protocol executes the ultimate SB 6 curtailment mandate: if the dispatched paired resource exhausts its duration, or if systemic scarcity triggers are met, the Telemetry Offsetting logic drops the shield, re-exposing the flexible portion of the load to SCED.

- **Safe Transition to SB 6-Compliant Backup Generation:** SB 6 encourages large facilities to utilize onsite backup generation to ensure operational resilience during grid emergencies. However, forcing a multi-hundred-megawatt data center into an uncoordinated, fast-acting curtailment can create voltage transients and hardware damage. The Telemetry Offsetting mechanism provides an operational bridge to achieve this statutory goal. By utilizing the paired resource as a "telemetry shield" during early scarcity indicators, the facility buys a manageable operational runway that is privately coordinated. This allows the load to start its SB 6-compliant onsite backup generation or shift IT load out of the system in line with the goals of SB 6.

[6] See Tex. Util. Code § 39.169, establishing PUCT study requirements, approval timelines, and potential hold-harmless conditions for BTM co-location arrangements with existing generation.
[7] See generally Tex. Util. Code § 37.0561, governing the interconnection of large flexible loads and establishing statutory authority to require load curtailment or the coordinated deployment of onsite backup generation during Energy Emergency Alerts (EEAs) to protect macro-grid reliability.

- **Auditable PUCT & ERCOT Oversight:** A core component of recent PUCT regulatory expansion is the ability to audit and penalize non-compliant loads. By mandating that the PCLR and the paired mitigating resource maintain strictly separated Resource IDs and distinct telemetry streams—regardless of whether they are interconnected FTM or within a PUN—ERCOT and the PUCT are granted unblended, auditable visibility. This ensures that the mitigating resource is never double counted toward multiple PCLR loads, ensuring provable regulatory compliance through ERCOT's native settlement system.

Ancillary Services Equality & Ratepayer Neutrality

Senate Bill 6 evaluates ratepayer protection through transmission infrastructure cost allocations (4CP) and stranded asset liabilities (CIAC/financial security). It does not classify native frequency integration or Ancillary Service (AS) impacts as an unlawful cost-shift violation.

The Telemetry Offsetting architecture natively aligns with this statutory boundary while requiring zero modifications to the *ERCOT Methodologies for Determining Ancillary Service Requirements:*

- **Mathematical Integration:** When the telemetry shield is raised, the PCLR operates identically to a firm load with zero dispatchable SCED range. Any uninstructed load fluctuations enter ERCOT's Area Control Error (ACE) logs as standard net-load forecast error.
- **Backward-Looking Statistical Sizing:** ERCOT sizes Regulation and Reserve products using backward-looking 95th-percentile net-load forecast error models. The PCLR's organic volatility is natively ingested by ERCOT's existing historical statistical models, ensuring future AS procurement automatically scales to account for the facility, the same as for firm loads on the system.
- **Equitable Treatment Across Large Loads:** By treating a shielded PCLR identically to an unconstrained firm Large Load during non-congested periods, Telemetry Offsetting establishes consistent market alignment. In accordance with Senate Bill 6's framework for equitable treatment across all Large Loads (≥ 75 MW), this construct holds PCLRs to the exact same physical and financial Ancillary Service standards as their firm load peers, keeping Texas ratepayers fully protected while maintaining a consistent operational standard across asset classes.

11. Big Picture: Positive-Sum through N-0 Neutrality & Scarcity Support

Telemetry Offsetting creates a two-state macro-economic paradigm designed to protect ERCOT operations across two distinct physical profiles, acting as a regulatory mechanism to incentivize new, fully dispatchable generation.

The Baseline State (N-0 Neutrality)

Under standard operating conditions, the Telemetry Offset PCLR and its paired mitigating resource are equivalent. Because the load's consumption is structurally offset by the localized injection of its paired resource, the addition of shielded load capacity does not drain the macro-grid's baseline reserve margin, exacerbate N-0 thermal creep, or accelerate the frequency of system-wide scarcity events.

The Scarcity State (N-1 Events & Net-Positive Relief)

During severe transmission contingencies (N-1) or system-wide Energy Emergency Alert (EEA) events, the physical carrying capacity of the grid drops rapidly. Rather than remaining islanded and neutral during these events—as is common with traditional Behind-The-Meter (BTM) configurations—the Telemetry Offsetting standard enforces a deterministically verified System Yield. The facility's load is curtailed, intentionally volunteering the unencumbered generation to serve the macro-grid. When the telemetry shield is dissolved, the facility's expanded

dispatchable range is natively ranked by its shift factor on ERCOT's Energy Management System (EMS), providing operators a highly visible relief valve to safely deploy the necessary physical counter-flow.[8]

12. Positive Externalities: A System-Level Assessment

The value of Telemetry Offsetting extends well beyond the protection of an individual facility. At scale, aligning the commercial imperatives of large load developers with the operational needs of the macro-grid delivers systemic benefits across four general categories outlined here:

- **Grid Visibility & Operations:** By providing a transparent commercial pathway for resources to remain on the grid, the architecture naturally incentivizes developers to avoid building off-grid or unsynchronized generation, preserving SCED optimization integrity. Furthermore, by ensuring the paired resource's dedicated capacity is accurately reflected in its COP, it grants the DRUC engine a clear picture of the node's net-neutral reality, preventing the exacerbation of out-of-market commitments.

- **Ratepayer Relief & Grid Utilization:** Existing grid infrastructure can often safely accommodate additional demand if managed carefully. Telemetry Offsetting maximizes the efficient use of available nodal headroom, delivering continuous value to ratepayers without requiring new subsidies. Additionally, allowing mitigated loads to energize safely ahead of permanent transmission build-outs mitigates long-term investment uncertainty for TSPs.[9]

- **Incentivizing New Dispatchable Generation:** The architecture leverages private technology-sector capital to construct new utility-scale batteries and gas generation directly on the grid. To ensure this framework drives net-additive generation to SCED rather than cannibalizing existing resources, it enforces temporal additionality via a strict "12-Month Rule". While these assets are privately financed to mitigate localized curtailment risk, they operate strictly within ERCOT's market systems; during system emergencies, they transition into an active scarcity reserve, providing operators with a highly visible, SCED-dispatchable layer of capacity to serve all Texans.[10]

- **Project Bankability & Capital Deployment:** Unmitigated curtailments present significant hurdles for traditional project finance. Telemetry Offsetting transforms this unpredictable exposure into a deterministic, contractually bounded operational profile. By providing a predictable framework for physical transitions, it creates the security institutional underwriters and lenders require, accelerating capital deployment into the ERCOT market.[11]

---

[8] See Arushi Sharma Frank, "Aggregations and data centers: If a resource shows up when the grid is straining, make it count," Utility Dive (July 31, 2025), arguing that utilities must develop integrated planning models that support the co-location of large-scale storage with load.

[9] See, e.g., Logan Goldie-Scot & Scott Gosselink, "Meeting load growth with clean, flexible power," Generate Capital (Sept. 2025), arguing that flexible data center demand could unlock up to 100 GW of new load nationally without requiring billions in immediate transmission upgrades, particularly in curtailed markets like ERCOT.

[10] Cf. Norris et al., "Data Centers and Generation Capacity over the Next Decade," Duke University Nicholas Institute for Energy, Environment & Sustainability (2026), modeling how flexible large loads and paired capacity reduce system peak stress, cut long-term capital expenditures, and lower overall retail power costs.

[11] See Arushi Sharma Frank, "The Texas Energy Reference Design: Stress-Testing Load Growth Challenges for an AI Century," Center for Strategic and International Studies (April 7, 2026), analyzing ERCOT's 410 GW interconnection queue and identifying Texas as the critical national laboratory for integrating massive, flexible data center loads with new generation.

Appendix: Telemetry Offsetting NPRR Protocol Mapping Summary[12]

This Appendix provides a summary of the Nodal Protocol Revision Request (NPRR) for the Telemetry Offsetting construct organized into five thematic groups spanning Sections 2, 3, 4, 8, 16, and a new Section 23 form of the ERCOT Protocols.

1. New Market Definitions (Section 2)

The NPRR establishes the foundational vocabulary required to link Telemetry Offset assets:

New Term	Definition
Telemetry Offset Resource Pair	A Provisional Controllable Load Resource (PCLR) and an offsetting Generation or Storage resource managed by the same QSE, sharing a Resource Node but having separate resource IDs.
Available Telemetry Offset Capacity	The lesser of: (a) the maximum MW quantity submitted on the paired Resource's Energy Offer Curve or Energy Bid/Offer Curve, as applicable; or (b) the paired Resource's telemetered High Sustained Limit (HSL).
Baseline LPC	The PCLR's underlying Low Power Consumption limit from its interconnection study, exclusive of any temporary shield-based increase.
Telemetry Offset Mitigated Offer Cap	The maximum SCED Step 2 offer price for a paired resource — the reference Adjusted Bid Cap for an ESR, or the lesser of the reference Adjusted Bid Cap or MOC for a Generation Resource.
Offset LPC MW	The amount, in megawatts, by which the telemetered Low Power Consumption (LPC) limit of a Provisional Controllable Load Resource (PCLR) operating as part of a Telemetry Offset Resource Pair exceeds its Baseline LPC.

2. Operations and Telemetry (Section 3)

Establishes the real-time physical parameters and Day-Ahead visibility rules:

- **The Telemetry Shield:** Authorizes the QSE to dynamically raise the PCLR's Low Power Consumption (LPC) limit up to the paired resource's Available Telemetry Offset Capacity.
- **Deterministic Scarcity Yield:** QSE drops the telemetry shield within 60 seconds of ERCOT declaring an Energy Emergency Alert Level 1-3, or if the paired resource suffers a derate/outage.
- **COP Alignment:** Mandates distinct Current Operating Plan (COP) reporting to prevent the Day-Ahead Reliability Unit Commitment (DRUC) engine from double-counting the shielding asset.

3. Pricing, Operations and SCED Mitigation (Section 4)

Establishes the automated economic constraints applied to the mitigating resource during Step 2 of SCED to ensure economic equivalency:

- **Telemetry Offset Mitigated Offer Cap:** Subjects Telemetry Offset Generator or ESR to the same Adjusted Bid Cap (ABC) methodology as the paired PCLR during non-competitive constraints.

[12] Disclaimer: The protocol mappings, definition structures, and operational parameters summarized below represent draft language prepared for stakeholder consideration. Final protocol drafting will be refined through the formal ERCOT Nodal Protocol Revision Request (NPRR) stakeholder review process and remains subject to modification prior to official adoption.

- **Lesser-than Generator Capping:** Generation Resources are capped at the lesser of the cost-based Mitigated Offer Cap (MOC) and the nodal ABC.

4. Deviation Exemptions and Settlement (Sections 6 & 8)

Shifts the burden of proof from real-time operations to post-Operating Day settlement:

- **Waiver of Penalties:** The waiver conditions are established in new Section 8.1.1.5; Section 6's existing deviation-charge provisions are referenced, not revised.
- **Volumetric Equivalency:** Requires the Available Telemetry Offset Capacity of the paired resource to be equal to or greater than the Offset LPC MW for every individual SCED interval within the 15-minute Settlement Interval. Because both resources share the same Resource Node and Settlement Interval, and because the paired resource's SCED offer is already capped at or below the PCLR's Adjusted Bid Cap (see Item 3 above), no RTLMP or other price calculation is needed.
- **Native Compliance Verification:** ERCOT performs the volumetric equivalency check, the EEA-latency cross-check, and the offer-sufficiency check directly from the telemetry, meter, and offer data it already receives under Section 10, Metering, and Section 11, Data Acquisition and Aggregation — consistent with the Telemetry Validation methodology ERCOT already applies to Aggregate Load Resources (ALRs).
- **Base Point Deviation Grace Period:** If the paired resource is unavailable unexpectedly while actively shielding, the QSE must lower the LPC limit, unshielding the PCLR immediately. The PCLR then has a grace period to ramp down to the firm base-line based on its Emergency Ramp Rate Down (EDRR) requirements. Any deviations during this period and the risk of having PCLR status revoked are waived.

5. Registration and Testing (Section 8, 16 & 23, Form AA)

Creates the guardrails for participation:

- **Commercial Operations Date (COD) Requirement:** Introduces a sworn Attestation (Form AA) requiring the mitigating resource to have achieved its Commercial Operations Date (COD) no more than 12 months prior to the PCLR's Initial Synchronization date, ensuring the framework drives new generation.
- **Paired Testing:** Recognizing testing obligations as a combined, coordinated pair to allow inflexible PCLRs to energize and follow SCED through the coordination of their paired resource. If the paired resource is unavailable the EDRR of the load is used.