

Black text = Submitted Question / Comment

Blue text = PJM Response

1. Compliance demonstration framework

The draft Attachment D-2 sets performance requirements but doesn't say how an LCL shows it meets them. Several provisions depend on study results (for example, the 2-second recovery default that applies "unless studies demonstrate that a slower recovery is necessary.") Other regions have addressed this directly: ERCOT uses its Model Quality Test framework, and MISO's draft requires a validated positive-sequence dynamic model. We recommend that PJM:

- Indicate where and when the compliance demonstration requirements will be defined;
 - LCLs will be required to demonstrate compliance with the ride-through and post-fault active power recovery requirements through the submission of applicable models and supporting technical information. PJM intends to provide guidance regarding modeling requirements, study assumptions, and the technical information necessary to support compliance assessments. PJM's objective is to establish a clear and consistent basis for demonstrating compliance across all applicable LCLs while ensuring that performance assessments reflect the characteristics of the interconnected Bulk Electric System.
- confirm whether the demonstration standard will be uniform across PJM or left to each TO; and
 - PJM intends to establish a uniform compliance demonstration framework applicable throughout the PJM footprint. PJM's requirements represent the minimum criteria that all applicable LCLs must satisfy and are intended to provide a consistent baseline across PJM. Consistent with the proposed Tariff provisions, Transmission Owners may establish more stringent requirements where necessary to address local reliability needs and system conditions. . This approach is intended to provide a predictable and transparent compliance framework while accommodating location-specific reliability needs when necessary.
- consider accepting OEM type-test certification plus documented protection settings as the baseline demonstration, with dynamic model studies required only when PJM needs one for a specific study.
 - PJM's current expectation is that demonstration of compliance will require the submission of validated EMT and positive-sequence power system dynamic (PSPD) models, together with sufficient technical information to evaluate facility performance under applicable disturbance conditions. OEM type-test certification and documented protection settings may provide useful supporting information; however, PJM does not currently envision such information serving as a substitute for the models and studies necessary to assess facility performance on the PJM system. Dynamic models are necessary to evaluate facility performance under system-specific disturbance conditions and to assess potential impacts on Bulk Electric System reliability. PJM anticipates that model quality and validation requirements for LCLs will follow principles similar to those established for

generators under NERC MOD-026 and MOD-033, whereby submitted EMT and PSPD models are expected to reasonably represent actual facility behavior and may require validation against measured performance data. The post-fault active power recovery requirement, requiring restoration of active power within 2 seconds following fault clearance, shall be the default requirement applicable to all LCLs. A slower recovery profile may be permitted only if PJM determines, based on planning assessments performed in coordination with the applicable Transmission Owner, that the default 2-second recovery requirement would result in adverse impacts on Bulk Electric System reliability. In such cases, PJM may establish an alternative recovery profile necessary to maintain Bulk Electric System reliability.

2. LCLs with behind-the-meter BESS

Section D-2.3 requires an LCL to remain connected through the voltage conditions specified in Table 1 and to draw at its pre-disturbance level while voltage stays within the continuous operation range. For a facility with behind-the-meter storage, this raises several questions:

- The 150% current limit. PJM should clarify what the 150% is measured against, - is this 150% of the interconnection limit of the campus as a whole, or of the computational component only (separating cooling load)? It would also be helpful to clarify what "maximum current drawn during normal operating conditions" refers to: pre-disturbance current, contract demand, nameplate, or a historical peak.
 - "Maximum current drawn during normal operating conditions" refers to the maximum steady-state operating current immediately preceding the disturbance. This limitation is intended to prevent an LCL from becoming a large current sink during disturbance recovery and exacerbating voltage recovery challenges on the Bulk Electric System. The requirement is intended to support predictable disturbance recovery behavior and limit adverse impacts on voltage recovery following fault clearing.
- Overvoltage events. If a co-located BESS absorbs power during a 1.10–1.20 p.u. event and pushes withdrawal above the pre-disturbance level, is that a violation? We suggest a tolerance band or ceiling instead of a fixed setpoint.
 - PJM is evaluating whether additional clarification is needed regarding facility behavior during high-voltage conditions, including the interaction of behind-the-meter energy storage systems and the applicability of the pre-disturbance load requirement during such events. . Any additional clarification would be intended to ensure consistent treatment of aggregate facility behavior while preserving the reliability objectives of the disturbance performance requirements.
- What LCL refers to. Several of these questions are dependent on the classification of 'LCL', and whether it's considered the total campus load or campus IT load only; with the proposed ride-through exception for cooling systems, clarification on what constitutes LCL would be appreciated.
 - An LCL consists of the computational load and the associated UPS, cooling, power delivery, and auxiliary loads necessary to support continuous operation of the facility. For purposes of evaluating compliance with the ride-through requirements, the LCL is evaluated as an integrated facility rather than as individual internal load components.

- "Transfer to backup generation." Does this include a co-located BESS or UPS that supports the LCL while it stays connected at the POD/POI? Or does it cover only schemes that physically island the LCL from the Transmission System?
 - PJM's objective is for the LCL itself to remain connected and ride through applicable disturbances. PJM has proposed clarifying language in M14B Attachment D-2 to make clear that compliance is expected to be achieved through ride-through of the Large Computational Load itself. A scheme that disconnects, transfers, substitutes, or otherwise replaces the electrical demand of the LCL with another resource during a disturbance would not, by itself, satisfy the ride-through requirement. The underlying objective is to avoid the transfer or substitution of load in a manner that masks the actual disturbance response of the LCL from the Bulk Electric System. Alternative solutions may be considered where ride-through of the LCL is demonstrated to be technically infeasible.

3. Cooling load carve-out (0.35 p.u.)

We support treating cooling loads separately. However, commercial chiller VFDs are typically configured to IEEE 1566 which permits shutdown for voltages below 0.65 pu, sustained beyond 500ms [1] [opencompute.org]. That means the cooling load may drop out inside the mandatory ride-through region. It would be helpful for PJM to:

- explain the technical basis for the 0.35 p.u. threshold, and consider whether a higher threshold better reflects how cooling equipment actually behaves; and
 - The basis for the 0.35 pu threshold is ERCOT's voltage ride-through requirement, which has undergone stakeholder review and approval. PJM selected this threshold to align with an established industry precedent while providing a practical distinction between critical computational load and supporting cooling equipment during severe voltage depression events. The threshold establishes a practical balance between recognizing the limitations of cooling equipment and maintaining continuity of the primary computational load during severe voltage disturbances.
- clarify what counts as "cooling load." Is it only chiller compressor drives, or does it also include CRAH/CRAC units, chilled and condenser water pumps, cooling tower fans, and their controls?
 - Cooling load is electrical load required to maintain the operating temperature, humidity, and environmental conditions necessary for continuous operation of the computational equipment within a Large Computational Load facility. Cooling load includes HVAC systems, chillers, chilled-water pumps, condenser-water pumps, cooling towers, CRAH/CRAC units, and associated control systems.
- Reflect equipment capability in the threshold. Firmware updates may not close the gap between 0.35 pu and where cooling equipment actually rides through.
 - PJM recognizes that the ride-through capability of commercially available cooling equipment may vary among manufacturers and facility designs. The proposed cooling-load exemption is intended to acknowledge these practical limitations while preserving ride-through expectations for the primary computational load. The proposed threshold is intended to provide a practical and technology-neutral framework that can be applied consistently across a range of facility designs and equipment configurations.

- Clarify how PJM intends to verify this requirement, as it will need to be monitored downstream of the POI at the cooling plant level. What are telemetry expectations for proving compliance?
 - The ride-through requirement is generally evaluated based on the behavior of the entire LCL at the Delivery Point or Point of Interconnection rather than through separate telemetry for individual internal load components. Consistent with other performance-based requirements, compliance focuses on the observable aggregate performance of the LCL at the Delivery Point or Point of Interconnection rather than the operation of individual internal components. Compliance is expected to be demonstrated through submitted technical information, facility design documentation, and applicable models and studies. PJM continues to evaluate implementation details associated with compliance verification.
- Acknowledge that commercially available chiller plants, if tripped, require several minutes to restart and return to pre disturbance state; this may influence how soon the POI returns to pre disturbance load level.
 - PJM recognizes that cooling equipment may require significantly longer recovery times than computational equipment following a trip. This consideration was one factor in proposing separate treatment for cooling loads within the ride-through framework. Accordingly, recovery of the total facility load may occur over a longer timeframe than recovery of the computational load alone.
- Consider whether there are aggregate system effects of exempting cooling loads; depending on the scope of LCLs in a given area, their cooling loads could still represent a substantial overall portion of electrical demand.
 - PJM recognizes that cooling load may represent a significant portion of total facility demand. The proposed exemption reflects a balance between practical equipment capabilities and reliability objectives. PJM will continue to evaluate the potential aggregate system impacts of cooling-load exemptions as stakeholder discussions progress and operating experience dictates. PJM expects the aggregate impact of cooling-load exemptions to be considered as part of ongoing planning assessments and future stakeholder review as LCL adoption and system experience evolve.

4. Consistency across Transmission Owners

Under the proposed Tariff §28.2, TOs may set requirements stricter than PJM's minimum. We recognize that local conditions sometimes call for this. Still, we urge PJM to make sure the result isn't a patchwork of TO-specific requirements, which would make design and procurement harder for LCLs across the footprint. We'd welcome guidance on how TO requirements will be reviewed and published. We also encourage PJM to align its requirements with NERC's large-load ride-through standards now in development, to limit duplicative efforts once those take effect.

- PJM's proposal establishes minimum ride-through requirements applicable throughout the PJM footprint. These requirements represent the baseline criteria that all applicable LCLs must satisfy, regardless of location. Consistent with the proposed Tariff provisions, Transmission Owners may establish more stringent requirements where necessary to address local reliability needs and system conditions. PJM recognizes the value of consistency across the footprint and will

continue to consider stakeholder feedback regarding transparency, coordination, and communication of any Transmission Owner-specific requirements. PJM also intends to monitor the development of NERC's large-load ride-through standards and consider alignment where appropriate.

5. Consequences for non-compliance

The deck notes that PJM will clarify the consequences of non-compliance. We'd appreciate more detail on PJM's current thinking, including how it will handle cure periods and equipment-driven limitations. We would be glad to work with PJM staff on this and to share how our facilities are designed.

- PJM is continuing to evaluate the appropriate compliance and enforcement framework associated with the proposed ride-through requirements. Additional details regarding non-compliance, corrective actions, cure periods, and implementation considerations will be developed.

6. Additional clarifying questions

- 2-second / 90% recovery. What is the technical basis for this default, and who performs the studies that could justify a slower recovery?
 - The 2-second recovery requirement is generally aligned with requirements adopted or proposed by ERCOT and MISO and reflects PJM's objective of minimizing prolonged post-disturbance load recovery effects that could adversely impact Bulk Electric System reliability. The requirement is intended to promote predictable post-disturbance recovery behavior and reduce the risk of large aggregate load restoration occurring over an extended period following fault clearing. Studies supporting an alternative recovery profile would be performed as part of PJM's planning assessment process in coordination with the applicable Transmission Owner.
- Six sags or swells within 90 seconds. How was this value chosen, and how does it compare with the ERCOT and NERC LLWG approaches?
 - The requirement was developed based on PJM operating experience with transmission line fault and reclose operations, which can result in multiple voltage sags and swells occurring over a relatively short period. The intent is to ensure that LCLs remain connected through repetitive disturbance sequences and do not repeatedly transfer load during transmission system restoration activities. The criterion is intended to represent credible transmission disturbance and restoration scenarios and provide assurance that LCLs can withstand multiple sequential disturbances without repeated load transfer.
- Frequency ride-through. Is the "should" in the 10% active-power band intended to be mandatory ("shall")? Also, are LCLs expected to participate in UFLS, and how does the UFLS exemption interact with the 57.0–58.8 Hz ride-through band?
 - The "should" in the 10% active-power band is intended to be mandatory and will be revised for consistency in future language updates. PJM specifies only the percentage of load that must participate in UFLS within each Transmission Owner area. Whether and how an LCL participates in UFLS should be coordinated with the applicable Transmission Owner. Ride-through requirements and UFLS requirements serve different reliability objectives, with ride-through intended to prevent unnecessary load disconnection during disturbances and

UFLS intended to support system frequency recovery during severe under-frequency events.

- LCL definition. Is there a quantitative threshold for "substantially supplied through power electronic conversion equipment"?
 - PJM's initially proposed definition used the phrase "substantially supplied through power electronic conversion equipment." PJM evaluated stakeholder feedback and changed to the following:

“Large Computational Load shall mean load facilities, including data centers and other computational facilities, with aggregate facility demand of 50 MW or greater and containing 25 MW or greater of maximum load served through power electronic conversion equipment, and meeting one or more of the applicability criteria below:”
- Alignment with other envelopes. Where PJM's envelope differs from MISO's draft (proportional reduction and mandatory recovery), the NERC LLWG SAR, or IEEE 2800, please explain the reasoning for the difference.
 - PJM's proposed disturbance performance requirements were informed by existing ERCOT requirements, recommendations developed through the NERC Large Load Working Group effort, and PJM's review of historical load-transfer events. Where PJM's proposal differs from other frameworks, those differences reflect PJM-specific reliability concerns and operating experience. In particular, the proposed requirements are intended to mitigate the reliability impacts of large-scale load transfer behavior and promote predictable disturbance performance during system recovery.