Stakeholder Feedback and IESO Response to Technical Requirements Document for Large Computational Loads Connecting to the Ontario Power System

Connection Assessments Department, IESO

Connecting Today.
Powering Tomorrow.

Table of Contents

1. BBA Consultants

BBA Consultants ("BBA") welcomes the opportunity to comment on the IESO's draft Technical Requirements for Large Computational Loads dated May 1, 2026. We commend the IESO for moving early to provide clarity to the sector - ahead of the NERC Large Loads Task Force's final guidance, and in parallel with other similar development across the industry including the ERCOT Large Load Working Group, AESO Phase I Large Load Integration, and others. The July 10, 2024 Eastern Interconnection event, in which a 230 kV fault cleared in 42-66 ms tripped ≈ 1.5 GW of voltage-sensitive data centre load across multiple sites and pushed system voltage to 1.07 p.u., underscores why this work is timely for Ontario.

BBA is a multidisciplinary engineering firm with deep specialization in power system planning, Ontario IESO interconnections, and the design of large industrial and data-centre/ computational load facilities.

General Observations

BBA supports the direction of the draft, which is well aligned with IEEE Std 2800-2022, and with discussions occurring in other ISO jurisdictions. Three thematic viewpoints frame our detailed comments that follow.

1. Constructability and CAPEX of the load-side ride-through requirements. Several clauses

- most notably the proportional active power reduction between 0.5 and 0.8 p.u. (Section 6.7.1(2)) and the 1-second active power restoration default - describe a load behaviour not natively achievable by today's commercially deployed double-conversion UPS fleet, which is designed to IEC 62040-3 and ITI/CBEMA standards. As written, hyperscale-sized data centre compliance would require supplemental grid-forming BESS, controllable load banks, or dynamic

VAR devices sized at hundreds of MW per site. BBA recommends pairing the requirements with (i) a defined **transition period** plus IESO-led vendor consultation with UPS equipment manufacturers and OEMs, and (ii) an explicit alternative compliance path - analogous to AESO Phase I - that accepts equivalent system-level outcomes via alternative mitigations that limits coincident load loss to a pre-defined system-acceptable value, or one defined by the outcomes of a system impact assessment.

IESO Response: We recognize that the ITIC curve is typically used to protect voltage-sensitive equipment within large load facilities during voltage disturbances. However, it is not intended to maximize the ability of these facilities to remain connected to the grid during such events. Because the ITIC curve is not aligned with our system's disturbance profile and timing, dedicated voltage ride-through curves designed for large loads to support grid reliability are required to ensure that connecting facilities can ride through power system faults.

Based on discussions with developers, OEMs, utilities, and consultants, we have not identified any technical or implementation challenges associated with the two aspects of the proposed requirements. With respect to recovery time, stakeholders have indicated that power restoration can occur almost instantaneously. The 1-second delay is therefore introduced as a margin to confirm that system voltage recovery is stable and sustained before reconnecting the load.

Overall, these requirements have been developed based on system needs rather than existing equipment standards. Equipment standards should evolve to meet system requirements, not the other way around, particularly when no technical barriers have been identified.

That said, we note that ERCOT, which has similar requirements, adopted a transition-period approach consistent with your proposal. Based on the experience and implementation process in ERCOT, mandatory compliance could begin on January 1, 2028. Given that few LCL projects are expected to enter in-service before that date, we do not believe it is necessary to establish an interim exemption period for these two technical aspects.

2. Vendor capability gaps in protective relays and UPS controllers. Several requirements assume settings and measurement windows not configurable in standard relays and UPS controllers in service today - notably the 0.5-s RoCoF averaging window (6.7.3), the 3–6 cycle frequency window (6.7.2), and the 25° phase-angle-jump capability (6.7.4).

Confirmation with relay vendors (e.g. SEL, GE Multilin, ABB, Siemens) and UPS OEMs (e.g. Vertiv, Eaton, Schneider, ABB, Toshiba, Mitsubishi, etc.) ahead of finalization would protect against requirements that are technically prescriptive but commercially unimplementable. Grid-interactive UPS product class is now becoming commercially available and supports programmable grid services; however, industry adoption is still gaining ground and availability is varied. Framing requirements in the transition period around current vendor capability would greatly benefit the data centre industry.

IESO Response: As noted above, these requirements have been developed based on system needs rather than existing equipment practices.

We do not agree with the characterization that these requirements are technically prescriptive but commercially unimplementable. While such capabilities may not yet be widely available in the market, there are no fundamental barriers to their implementation. Where there is a clear system reliability need, the market and technology will evolve to provide commercially viable solutions.

It is important to note that similar requirements have already been applied to inverter-based resources, which likewise evolved from being commercially unavailable to becoming standard features over time.

3. Differentiation by connection voltage and project size. In some cases, the draft applies one set of thresholds (e.g. 20 MW/min ramp) across all projects ≥ 10 MW, whereas a tiered classification system or zonal-relevant values may be more appropriate. BBA recommends the IESO consider a tiered structure indexed to POI voltage (115 / 230 / 500 kV) and project size, so that a 20

MW distribution-connected facility is not subject to the same technical burden as a 500 MW campus connected at 500 kV.

IESO Response: The primary purpose of the load ramping requirement is to prevent depletion of the AGC fleet. This is a system-wide concern and is independent of connection location or voltage level. Therefore, it is not appropriate to differentiate requirements based on connection voltage or project size.

4. Equivalence with other industrial load facilities. The guiding principle should be that large computational load requirements are treated consistently with other large industrial loads - neither subject to more lenient standards nor disproportionately stringent requirements.

IESO Response: We are aligned with NERC's direction and are prioritizing large computational loads. The rationale is that most of the large loads we are currently dealing with fall into this category. These loads introduce the greatest reliability risks and challenges, which need to be addressed urgently. We will subsequently address other industrial load facilities and ultimately establish requirements applicable to all large load facilities.

The new requirements are intended to address specific load behaviors exhibited by these facilities. They will not be valid to facilities that do not exhibit such behaviors. For example, regarding load fluctuation requirements, if a facility does not demonstrate load fluctuations, then even if such requirements are specified in the SIA, they would not be applicable or meaningful for that load.

Detailed Observations

The following comments provide our detailed perspective to various technical clauses proposed in the

Technical Requirements for Large Computational Loads Connecting to the Ontario Power System V1.0.

1. Section 5 - "All primary equipment data." "Primary" is undefined and may be ambiguous to some readers. Recommend specifying whether this is referring to primary connection voltage (high voltage) or referring to main ("primary") power system equipment.

IESO Response: It is a power system term referring to high-voltage equipment. We will clarify it in the technical document to eliminate any ambiguity.

2. Section 5 - Type of large computational loads. The current taxonomy (cloud, AI training, AI inference, crypto) describes use cases. Grid impact, however, is driven by hardware class (GPU clusters, ASIC, CPU-dominant, liquid- vs. air-cooled). Recommend the data submission also list, or primarily focus on, aggregate hardware type, which drives the fluctuation and ramp profile and varies most rapidly with technology generation.

IESO Response: This is a valid point. Instead of replacing the type, we will add the hardware class and collect both.

3. Section 5 - "Post-fault active power recovery (PFAPC) time." It is unclear what the IESO would expect to be listed against this parameter if it is proposing to establish 1 second as the restoration time after system disturbance (ref. 6.7.1) (also discussed further below).

IESO Response: Yes. We will make the wording consistent.

4. Section 5 - "UDM/Generic composite model for PSS/E and DSA tools." We recommend a joint technical working group – along with the IESO – to develop and publish a guiding technical specification (PSS/E version, DSAT version, parameter ranges, benchmarking criteria vs. EMT modeling, and acceptance tests) analogous to current IBR modelling guidance. PSCAD/UDM models capturing data-centre load behaviour are not generally available from OEMs and will be built by the applicant; an Ontario IESO-led technical working group, publishing an acceptance template would be materially beneficial to the industry.

IESO Response: The IESO will consider this recommendation. Thanks for this offering.

5. Section 6.4 - Periodic Oscillations. Recommend defining "periodic" with a minimum duration threshold and an assessment window (e.g., oscillations sustained beyond X seconds/ minutes at the POI within the sub-synchronous band). As written, any transient oscillations during a workload step, which are inherent in many data centre use cases, would technically be noncompliant. Aligning Section 6.4 with the load-fluctuation assessment under A.7 closes the interpretation gap.

IESO Response: The IESO will specify a minimum duration threshold to ensure the requirement can be effectively implemented and enforced. Thanks for your recommendation. Yes, we mean that any transient oscillations are not allowed. They are inherent on the IT loads side but can be smoothed out by appropriate UPS control so that they don't show up at the POI.

6. Section 6.6 - Load Ramping. A homogeneous 20 MW/min cap is significantly limiting for large projects. We recommend either (i) scaling the ramp rate with project size and POI voltage, or (ii) making it explicitly negotiable per project through SIA, with 20 MW/min as the default for projects connected at lower system voltages.

IESO Response: This is the same comment as item (3) in General Observations.

7. Section 6.7.1(2) - Proportional load reduction between 0.5 and 0.8 p.u. This is, in our view, the single highest-impact clause in the draft requirements. Modern double-conversion UPS systems transfer the DC-bus to battery on a voltage sag, similar to droop response. They are activated in this range to protect downstream IT equipment per IEC 62040-3. Requiring, instead, a proportional modulation would require either supplemental BESS / grid-forming devices at hundreds of MW per site, or UPS firmware / topology changes not yet uniformly available.

BBA recommends the IESO that in the **transition period**: (i) re-frame this as a permitted MW ceiling on load reduction based on SIA outcomes rather than a mandated proportional response; (ii) consider a vendor-readiness assessment alongside the requirement; (iii) adopt a reasonable transition period with grandfathering for existing projects, or projects already in SIA queue; and (iv) accept alternative site-level control logic as equivalent compliance.

IESO Response: This is the same comment as item (1) in General Observations.

8. Section 6.7.1(2) - 1-second active power restoration default. As specified, this is significantly faster than typical industry-practice for load restoration, especially for critical loads such as data centres. It is common practice in the industry to wait a few minutes for power restoration to stabilize before load restoration back to utility power. Compliance would require site-wide supplemental BESS or controllable load banks. BBA recommends (i) revisiting the default time proposed with the recovery time discussions happening in the NERC LLTF and other forums and/ or explain the explicit need for a 1-second restoration; and (ii) requiring staged restoration across UPS blocks so that simultaneous reconnection does not itself create a voltage step at the POI, and also preserves the required data centre reliability and resilience philosophies.

IESO Response: This is the same comment as item (1) in General Observations.

9. Section 6.7.2 - Frequency Ride-Through measurement window. A 3–6 cycle window is not universally available in commercially deployed protective relays and controllers. We recommend (i) alignment with relay vendors (SEL, GE Multilin, ABB) and UPS OEMs prior to finalization; and (ii) a more flexible approach would be to rely on established measurement algorithm families (e.g., zero-crossing, DFT-based, PMU-style estimators).

IESO Response: Please refer to our response for item (2) in General Observations.

10. Section 6.16 - Commissioning Tests and Performance Validation. Please refer to our comments provided for Item #4.

Closing

BBA thanks the IESO for the opportunity to comment, and for the open and constructive engagement process to date. We would welcome the opportunity to participate in one-on-one technical sessions ahead of the June 18, 2026 public webinar, where we can offer detailed input from our active data-centre design engagements as well as our subject-matter-experts.

2. Enova Power Corporation

The proposed technical requirements appropriately recognize that large computational loads, such as data centres and AI-related facilities, may behave differently from traditional industrial loads and can introduce unique reliability and operational considerations to the Ontario power system.

From an LDC perspective, we generally support the IESO's proactive approach to establishing technical and operational requirements for these types of facilities, particularly in the areas of:

- load ramping limitations,
- ride-through requirements,
- restrictions on automatic reconnection,
- enhanced telemetry and monitoring,
- and customer responsibility for mitigation infrastructure and system impacts.

We also support the IESO's recognition that large computational loads may require enhanced operational coordination between the customer, transmitter, distributor, and the IESO.

Please see the following comments and questions organized by topic area below.

1. Clarification of Operational Roles and Responsibilities

Further clarity regarding operational authority and coordination between the IESO, transmitters, and LDCs would be beneficial, particularly with respect to:

- energization and restoration sequencing,
- curtailment implementation,
- reconnection approvals,
- and emergency operating procedures.

Additional guidance on how operational decision-making responsibilities will be coordinated for distribution-connected large computational loads would be helpful.

IESO Response: We expect the operational roles and responsibilities to align with those of existing load facilities above 10 MW, with the addition of new technical requirements specific to large computational loads.

2. Distribution Connection Process Alignment

Given that the requirements apply to distribution-connected projects above 10 MW, additional guidance regarding how these requirements integrate with existing distribution connection processes would be beneficial.

Consideration should be given to:

- standardized processes,
- study coordination timelines,
- technical review responsibilities,
- and communication protocols between the IESO, transmitters, and distributors.

IESO Response: The process will remain the same as for existing load facilities (>10 MW). New requirements will be implemented through the existing SIA process, rather than the distribution connection process. Currently, there are no additional requirements for large computational facilities below 10 MW.

3. Staged Energization and Restoration Practices

The proposed ramp-rate limitations and ride-through requirements are supported. However, consideration should also be given to developing:

- standardized staged energization practices,
- restoration sequencing expectations,
- and operational guidelines for large computational load facilities.

These practices could help reduce operational risk during commissioning, restoration, and abnormal system events.

IESO Response: The process will remain the same as for existing load facilities (>10 MW).

4. Mitigation Infrastructure and Cost Responsibility

The document refers to the potential need for:

- SVCs,
- STATCOMs,
- capacitor banks,
- grid-forming batteries,
- Remedial Action Schemes (RAS),
- and other mitigation infrastructure.

Additional clarity regarding:

- cost responsibility,
- ownership,
- operation,
- and long-term maintenance obligations would be beneficial for both distributors and customers.

IESO Response: The implementation of these requirements will remain consistent with that for existing load facilities (>10 MW). The Local Distribution Company (LDC) will act as the connection applicant for the distribution connected project and will be responsible for implementing the SIA requirements. The LDC will determine the allocation of cost responsibility, as well as ownership, operation, and maintenance arrangements between itself and the customers.

5. Telemetry and Monitoring Requirements

The proposed telemetry, synchrophasor, and dynamic disturbance recording requirements are supported, as they improve system visibility and operational awareness.

Additional implementation guidance regarding:

- data ownership,
- communications requirements,
- interoperability expectations,
- and distributor operational access to telemetry data where applicable would help support consistent implementation across the province.

IESO Response: The PMU and DDR requirements are not applicable for the distribution connected projects. The IESO has not yet identified any additional telemetry. The telemetry and monitoring requirements generally remain the same as for existing load facilities (>10 MW). Additional telemetries, if applicable, will be identified during the SIA process.

6. Behind-the-Meter Generation and Energy Storage

Additional clarification regarding the treatment of:

- backup generation,

- battery energy storage systems (BESS),
- and other behind-the-meter resources would be beneficial.

Additional guidance regarding:

- anti-islanding expectations,
- synchronization requirements,
- operational visibility,
- and restoration coordination would support consistent implementation and operational coordination across the province.

IESO Response: Backup generation is typically not operated in parallel with the system. It is up to the load customers to determine how their loads are supplied. For BESS and other behind-the-meter resources, the IESO will assess them as part of the SIA and specify requirements for operational visibility and restoration coordination. However, anti-islanding expectations and synchronization requirements are outside the scope of this assessment and remain the responsibility of the LDCs.

7. Load Forecasting and Capacity Planning

Large computational loads may materially impact regional load forecasting, distribution planning, and transmission planning processes. Further guidance regarding:

- long-term load forecasting expectations,
- phased capacity allocation,
- and planning coordination between the IESO, transmitters, and LDCs would be valuable as these types of projects continue to grow within Ontario.

IESO Response: This is a valid topic; however, it falls outside the scope of this stakeholder engagement. The IESO will engage with LDCs directly when addressing this subject.

8. Standardization of Technical Requirements

Given the complexity of the proposed modelling and technical submission requirements, consideration should be given to developing:

- standardized templates,
- model submission requirements,
- and technical acceptance criteria to improve consistency and reduce administrative burden for all parties involved.

IESO Response: Upon completion of this stakeholder engagement, the updated technical requirements will be incorporated into the Market Rules. Corresponding updates to model submission and performance validation requirements will be reflected in the relevant Market Manual(s).

Closing Comments

Overall, the proposed framework represents an important step toward addressing the operational and reliability impacts associated with large computational loads in Ontario.

We appreciate the IESO's proactive engagement on this topic and look forward to continued collaboration as the requirements evolve and implementation approaches are further refined.

3. Hydro Ottawa Limited

Hydro Ottawa is generally supportive of the proposed technical requirements for large computational loads and its objective to protect the reliability of the integrated power system.

The utility seeks clarification on some elements of the proposed requirements.

Hydro Ottawa seeks clarification on the following proposals:

1) Section 6.1 states that automatic load transfer between different supplies is not acceptable. Hydro Ottawa requests clarification on the intended scope of this restriction. Specifically, how will this requirement impact standard customer emergency backup arrangements (e.g., ATS/open-transition schemes)?

IESO Response: We will revise the wording to clarify the prohibition on automatic load transfer applies to load transfer between different POIs, not to the load transfer from the system supply to backup generation.

2) Section 6 assigns responsibility for power quality performance to the transmitter but does not clearly identify applicable requirements or responsibilities. Hydro Ottawa requests clarification regarding applicable standards. In addition, which parties (transmitter, distributor, or customer) are responsible for enforcing and executing mitigation measures for potential power quality impacts from large computational loads?

IESO Response: The document does not "*assign*" responsibility for power quality requirements to the transmitter. Rather, it states that power quality falls within the transmitter's accountability. Therefore, please contact the transmitter regarding any power quality related matters, including applicable standards and mitigation measures.

3) Section 6.14 specifies that the IESO may require additional telemetry quantities if they are deemed applicable to the project. To support proactive utility planning and customer guidance, Hydro Ottawa seeks clarification on the specific monitoring parameters expected to generate this data (e.g., UPS status, MW, backup generation).

IESO Response: The IESO has not yet identified any additional monitoring variables, except for indicating the need for monitoring tools such as PMUs and DDRs which are not applicable for distribution connected projects.

4) Section 6.7.1 (2) states "After the system disturbance is cleared and the POI voltage recovers within the ORTAC post-contingency voltage range, the project shall restore its active power

withdrawal from the system to the level that exists prior to the system disturbance within a restoration time that is set to 1 second by default and configurable by the IESO based on system assessments." Hydro Ottawa suggests that a 1 second recovery window is an aggressive approach, and therefore requests clarification as to whether restoration should be staged rather than introducing a large load to the system all at once. The necessity of a staged approach is dependent on the project size. If not, why not?

IESO Response: We don't consider the staged recovery is necessary. When a large amount of computational load transfers to UPS during a fault, it reduces its demand on the system. Once the fault is cleared, this can result in a temporary excess of generation over the load, which may impact system frequency and lead to local overvoltage conditions. The magnitude of this imbalance is directly related to the amount of load transferred to UPS. Therefore, fast restoration of the load without ramp-rate limitations is required.

That said, this recovery process will be modelled and assessed in our system studies. If any adverse impacts are identified, the SIA may require different recovery timelines for individual projects.

5) Finally, Hydro Ottawa seeks clarity on how the proposed changes would apply to existing >10MW customers who are currently midway through the connection process.

IESO Response: The new requirements are applicable for the projects which are still in the interconnection process and haven't been given conditional approval by the IESO. They are not applicable for those projects which have obtained conditional approval from the IESO. We will clarify it in our technical document.

4. Hydro One Network Inc.

Comment 1 from Michael Xavier by email

Section 6.1 Connection Arrangement : No internal single point of failure within the project results in a load loss above 600MW or a lower value identified by system assessments

The 600MW limit is as listed in ORTAC for load security. Load security is to ensure a reliable supply under contingencies and maintaining acceptable levels of interruption. The 600MW limit was a somewhat arbitrary construct based on 3 DESNs worth of load being an acceptable amount of load that could be lost. As such I would consider removing data centre load from load security requirements as the loss of this load would only impact one customer and not a broader load interruption which is the basis for the 600MW limit. I would suggest that restrictions on the amount of load interrupted be based on system impact.

IESO Response: Due to the potentially large magnitude of computational loads, we propose establishing a maximum allowable load loss following an internal fault at such facilities. A limit of 600 MW is consistent with the maximum allowable load loss associated with a system event in our system. This requirement is intended to mitigate broader reliability impacts that may result from an internal fault at a large computational load facility.

Although the 600 MW threshold is derived from the maximum load loss associated with the loss of two transmission elements based on load security criteria, it is important to note that this requirement is not intended for load security purposes. Rather, the objective of specifying a maximum load loss limit is to manage and reduce potential adverse impacts on overall system reliability. The rationale for this requirement is outlined below:

(1) Specifying a maximum load loss limit of 600 MW following an internal fault at a large computational load facility provides an effective means of safeguarding both the system reliability and operability. Such facilities represent highly concentrated and fast-responsive electrical demand. An uncontrolled internal fault could result in the sudden disconnection of a very large block of load, creating an immediate imbalance between generation and demand. This imbalance may lead to frequency excursions that could challenge generator governor response and automatic generation control (AGC).

(2) The limit also constrains voltage transients associated with abrupt load rejection. Large computational loads are typically dominated by power electronic equipment that is sensitive to voltage variations and contributes to dynamic reactive power behavior. A sudden loss of a large load block can lead to overvoltage conditions and reactive power imbalances, particularly in areas with limited voltage support.

(3) The 600 MW threshold is aligned with the IESO's contingency-based planning and operating framework, which is designed to withstand credible single contingencies without widespread disruption. By keeping the maximum load rejection within this scale, the impact of an internal fault remains consistent with disturbances that the system is designed to manage. This enables operators to rely on established procedures and controls, improves operational predictability, supports accurate contingency analysis, and reduces the likelihood of emergency interventions.

(4) An alternative approach, i.e. studying the loss of load on a case-by-case basis to identify the maximum allowable load loss for each project, may look reasonable during the interconnection process. However, this approach risks losing a consistent system-level perspective on the issue introduced by individual projects. If the loss of load at a given project becomes a limiting factor, it would need to be treated as a specific contingency in real-time operations. There is a significant risk that such contingencies may not be consistently captured or monitored, increasing the potential for cascading impacts.

(5) The IESO has evaluated various values for the maximum load loss limit and concluded that a 600 MW threshold provides a practical and technically justified balance. It accommodates the scale of modern computational facilities while ensuring that any internal

fault results in a disturbance that is controllable, predictable, and consistent with the reliability and operability requirements of the Ontario power system.

Comment 2 from Chester Li

In voltage ride-through, may point out the Vmax is RMS measurement based, and the per unit is not based on instantaneous peak. At mean time, may require ride-through of transient over-voltage up to 2 per unit as defined in Appendix 2 of TSC.

p.16, typo in RAS

IESO Response: System voltage is generally defined using RMS voltage in power industry as well as in our existing Market Rules and Market Manuals, unless otherwise specified. In addition, the per-unit (pu) values are not affected by whether RMS or instantaneous peak values are used. As such we don't see any confusion over here. That said, we will try to specify RMS when applicable.

The 2 per unit value specified in Appendix 2 of the TSC refers to equipment capability to withstand capacitor switching surges, not from system reliability perspective. It is not a protection setting that can be adjusted. In addition, as this requirement is already included in the TSC, there is no need to repeat it.

Thank you for pointing out the typo!

5. SLR Consulting

1. To confirm the process outlined by the IESO in the document, the following steps will occur:
    a. Proponent provides benchmarked PSSE model with UDM & composite model AND EMT model (with OEM EMT UPS)
    b. IESO performs review of the models and performs analysis.
    c. IESO recommends changes (i.e. BESS, STATCOM)
    d. Would capacity be reserved and ability to proceed with CCRA at this point (comfort letter, draft SIA)? Or would proponent need to provide updated benchmarked PSSE model with the updated EMT model including the BESS/STATCOM changes? It should be noted in most cases the proponent would be far from selecting a final UPS & dynamic equipment vendors, and primarily need capacity reserved prior to proceeding with significant costs.

IESO Response: Capacity is not reserved at the time a comfort letter and draft SIA are issued. Even if the IESO completes the study and issues a Notice of Conditional Approval, this does not constitute capacity reservation. Capacity is only reserved once the project becomes committed, as defined by the criteria in Market Manual 1.4.

The decision to proceed with the CCRA is at the discretion of the transmitter.

2. Has any consideration been given to adjusting the interconnection process to a staged review to provide proponents with certainty of connection capacity prior to significant spend on modeling efforts like other jurisdictions?

   For example, in Alberta DTS/STS are established via preliminary PSSE models with detailed engineering models (PSSE & PSCAD) submitted during additional stage gates prior to interconnection. This allows proponents to perform procurement activities and selections prior to significant modelling efforts & cost. The expectation of some of the data requested in #5 Project Data Requirement is detailed and requires selection of specific vendors to complete which is likely to result in SIA resubmissions when procurement inevitably changes product selection.

IESO Response: Yes, the IESO has adopted a flexible interconnection process to reflect the level of project information available at different stages. During the early stage, when detailed project design is not yet available, the IESO conducts steady-state studies only and issues a comfort letter along with a draft report to the applicant, indicating the available capacity for the project.

Based on this comfort letter, the applicant can proceed with procurement activities and vendor selection. Once specific vendors are selected, the applicant must provide detailed project design including vendor-specific models. The IESO will then complete the remaining studies and issue a Notice of Conditional Approval for the project.

3. It is currently difficult to obtain OEM EMT models (currently) for UPS systems as vendors are currently developing them with Manitoba Hydro (or other PSCAD model providers). Has the IESO considered adjusting its SIA processes to allow future submissions to use generic models first, followed by (i.e. during detailed design - after procurement activities but prior to energization) OEM models later? The question relates to #2 and confirming load availability prior to engaging in significant modelling effort.

IESO Response: Same as above for #2.

4. Data Centres typically transfer between grid & generators using automatic transfer switching (with breaker pairs or purpose-built Automatic Transfer Switches). While UPS/IT loads could be controlled to ramp during re-connection, mechanical loads would generally occur as switching events. The linear provision in 6.6 Load Ramping could be interpreted as not permitting this standard form of load transfer.

IESO Response: The IESO acknowledges that non-UPS mechanical and cooling loads are typically connected through switching events during reconnection and may not be able to comply with the linear ramping provision. Thank you for bringing this to our attention.

To make the requirement more practical, we have revised it by removing the linear ramping provision. Instead, we have added that the maximum allowed step change is 5 MW over any 15-second period. Under the revised requirement, each non-UPS mechanical or cooling equipment, which is generally less than 5 MW each, may be switched during reconnection. However, to avoid a large step change and ensure compliance with the revised ramping limit, these loads should not be switched simultaneously. They should be controlled such that no more than 5 MW of load is connected within any 15-second period. We believe this revision addresses your concern.

5. Will changes to BESS & UPS equipment selection impact capacity reservation? Given the early stage of the SIA submissions, it's likely most proponents will need the ability to change based on market (cost & availability) of equipment after initial SIA and CCRA payments.

IESO Response: No. It should be noted that capacity is reserved only after the project becomes committed. The conditions for a project to become committed are documented in Market Manual 1.4.

Once the project is committed, changes to BESS and UPS equipment selection will not impact on the reserved capacity; however, they may result in additional equipment upgrades within the project.

6. As part of A.9, is the IESO willing to discuss additional grid-interconnected behind-the-meter generation (gas turbines, etc.) at the proponent facility to solve possible load availability challenges (i.e. non-export)?

IESO Response: Yes, the IESO supports the behind-the-meter generation option at the proposed facility to address system capacity challenges. The IESO has established performance requirements for such generation and will conduct system impact assessments based on the operational philosophy, considering the presence of behind-the-meter generation.

6. Toronto Hydro-Electric System Ltd

Toronto Hydro supports the IESO's proactive efforts to establish a framework for large computational load (LCL) connections, recognizing the distinct operational characteristics of AI-driven data centres and similar facilities. Toronto Hydro's comments focus on enhancing implementation clarity, particularly with respect to roles and responsibilities, timelines, and coordination between distributors and the IESO.

Clarification of Roles and Responsibilities

Greater clarity on roles and responsibilities is needed to support consistent implementation and reduce uncertainty regarding accountability. Toronto Hydro notes that, while the proposal does not explicitly reassign responsibility for meeting reliability standards, it introduces expanded assessment expectations for distribution-connected LCLs. In doing so, the framework appears to extend elements of bulk system reliability assessment into the distribution connection process without clearly defining whether responsibility for conducting or validating these assessments remains with the IESO or is expected to be undertaken, in part, by distributors.

Clear direction on how the IESO, transmitters, distributors, and connection applicants will perform and validate reliability assessments – including demonstrating compliance with applicable standards (NERC, NPCC, ORTAC, TSC, and DSC) – would significantly improve implementation clarity.

IESO Response: The new requirements do not change the existing roles and responsibilities of the IESO, transmitters, LDCs, or connection applicants.

LDCs remain the connection applicants for distribution-connected LCLs. The new requirements will apply to LDCs, and it is up to each LDC to determine whether to pass these requirements down to LCL customers.

If you have limited experience with applying technical requirements for existing load facilities, you may refer to the established roles and responsibilities used when dealing with DERs.

Alignment of Timelines and Processes

Alignment between IESO and distributor processes is critical to support efficient project execution and maintain existing customer service levels. Additional transparency regarding timelines for IESO reviews, including SIA-related assessments, would improve coordination and planning.

As LDCs are subject to established connection obligations and timelines, integrating any new IESO requirements with existing distributor processes will help avoid uncertainty or unintended delays and support effective alignment of timelines across parties.

IESO Response: The above comments can be also applicable for timelines and processes.

Distributor Visibility into Transmission-Connected LCLs

Toronto Hydro notes that improved visibility into transmission-connected LCLs within or near its service territory will support effective planning and operations. These facilities can materially affect local planning, distribution infrastructure, and customer coordination. Maintaining distributor access to information on transmission-connected LCLs will ensure alignment between bulk and distribution system planning.

IESO Response: The new technical requirements introduced are not expected to change distributor visibility into transmission-connected projects. We do not see a particular need for additional focus on LCLs. The IESO and transmission owners are responsible for transmission-connected LCLs, including assessing their impact on system reliability and advising LDCs of any adverse effects during system planning, interconnection process, and operations.

Ongoing Engagement and Coordination

Toronto Hydro underscores that continued stakeholder engagement will support the practical implementation of the proposed modelling requirements. Ongoing collaboration will ensure that requirements remain implementable, rely on available data from customers and equipment manufacturers, follow appropriate sequencing, and align with existing connection processes – particularly given the rapid evolution of AI and data centre technologies.

As the framework evolves through Market Rules updates and related guidance, continued engagement will ensure that requirements remain clear, practical, and consistently applied across all parties.

IESO Response: The IESO agrees with Toronto Hydro on the importance of continued stakeholder engagement and coordination. Ongoing collaboration will be critical to ensuring that the requirements remain clear, practical, and consistently applied across all parties. Thanks!

7. XMight Smart Energy Corp.

XMight Smart Energy Corp. submits this voluntary technical input as a contributor to the Open Compute Project (OCP) Rack & Power and DCF LVDC work streams.

We note the framing in the IESO's engagement: data centres are not dispatchable resources; the reliability framework therefore rests on automated ride-through and physical recovery behaviour, rather than active operator dispatch. Our comments are technical, organised around clauses where the requirement language appears either internally inconsistent or — given typical AC-DC PSU design — operationally difficult to apply.

Where we reference the parallel ERCOT NOGRR 282 stakeholder process and the Data Center Coalition's formal comments on it, we do so to share lessons surfaced through that process rather than to recommend direct adoption of specific ERCOT language. References are at the end.

Editorial note — "PFAPC" in §5

§5 lists "Post-fault active power recovery (PFAPC) time (seconds) for UPS loads" as a required data element. The established industry term is PFAPR — Post-Fault Active Power Recovery (used in ERCOT NOGRR 282 §2.15(3) and in broader power systems literature). The §5 data requirement does not cross-reference the normative behaviour set out in §6.7.1(2). We suggest IESO either adopt PFAPR consistently in both §5 and §6.7.1(2), or remove the abbreviation in §5 and rely on the descriptive language in §6.7.1(2); either approach should add an explicit cross-reference from §5 to §6.7.1(2).

IESO Response: Thanks for pointing it out. We will remove the abbreviation in Section 5 to eliminate any confusion.

Issue 1 — §6.1 vs §6.7.1(2): Internal Contradiction on Backup Energy Engagement During Deep Sag

§6.1 states that "automatic load transfer between different supplies from system to the project is not acceptable." §6.7.1(2) states that during a voltage sag, the load "may cease to withdraw power from the system only when the POI voltage goes below 0.5 pu." These two clauses appear to be in tension.

The physical reality at V = 0.5 pu is that two things must happen simultaneously for the facility to meet the §6.7.1 ride-through requirement: (a) the load must current-limit at the PSU input — both to avoid drawing more current from an already-stressed grid and to protect the equipment; and (b) internal backup energy must engage to support the IT load while AC supply is below the PSU's operational range. Without (b), (a) cannot be sustained — the IT load loses its DC bus and the facility trips, which is the opposite of "ride-through."

Internal backup energy engaging automatically when AC supply falls outside the PSU's operational range is the architectural purpose of any backup energy device — facility UPS or rack-internal BBU. If §6.1 is read literally to prohibit this engagement, it prohibits the behaviour §6.7.1(2) requires.

Recommendation: Amend §6.1 to clarify that the prohibition on automatic load transfer applies to inter-supply switching between independent facility-level power supplies during normal operation (e.g., system supply vs. backup generation) and does not apply to automatic engagement of internal energy storage during a voltage disturbance within the ride-through bounds of §6.7.1.

IESO Response: Thanks for pointing it out. We will revise the wording to clarify the prohibition on automatic load transfer applies to load transfer between different POIs, not to the load transfer from the system supply to backup generation.

Issue 2 — §6.7.1(2) vs §6.7.2–§6.7.5: Concurrent Voltage and Frequency Events

§6.7.1(2) permits the load to cease withdrawing power when POI voltage falls below 0.5 pu. §6.7.2–§6.7.5 each state that the use of UPS is not allowed during the ride-through period for frequency, RoCoF, phase angle, and V/Hz events.

In a cascading disturbance, voltage and frequency excursions occur concurrently. A simultaneous V < 0.5 pu and f < 58.8 Hz condition triggers conflicting instructions: §6.7.1(2) permits backup engagement on the basis of voltage; §6.7.2 prohibits it on the basis of frequency. Without a defined priority, two operators applying §6.7 in good faith may implement opposite responses to the same event class.

Recommendation: Publish a priority resolution rule defining facility behaviour under simultaneous voltage + frequency / RoCoF / phase / V/Hz conditions, so operators and equipment vendors can design control schemes deterministically.

IESO Response: This is a valid point. Thank you for pointing it out! Switching to the UPS supply during voltage depression is permitted due to the equipment limitations of UPS loads. No such limitations apply under high voltage, frequency, RoCoF, phase jump, or V/Hz events. Therefore, in concurrent events, switching to the UPS supply is allowed whenever voltage depression occurs.

As such, we will revise the wording for frequency, RoCoF, phase jump, or V/Hz events, switching to UPS supply is not allowed during the ride-through period if there is no concurrent voltage depression.

Issue 3 — §6.6: Load Ramping — POI-Level vs Internal Workload Swing

§6.6 requires linear ramp rates not exceeding 20 MW per minute during normal operation. Generative AI workloads exhibit internal power swings of 30–80% of nameplate within sub-100 ms intervals (training-step boundaries, batch starts, model loads). At the facility level — as observed at the POI — these internal swings are absorbed by PSU bulk capacitance and rack internal energy storage; POI-observable ramp is substantially smoother than rack-internal behaviour.

Recommendation: Clarify that §6.6 applies at the POI, not at the internal compute-load level, and that internal compute power swings absorbed within the facility power architecture are not subject to the 20 MW/min constraint.

IESO Response: This is a valid point. The load ramping requirement applies at the project's Point of Interconnection (POI) rather than to internal load variations within the facility.

We will revise the wording to clarify that the load ramping constraint is applied at the project POI.

Issue 4 — §6.7.1(2): Proportional Reduction and 1-Second Restoration — Phased Path Consistent with ERCOT / DCC Experience

§6.7.1(2) requires the load to (i) reduce active power proportionally to voltage during a 0.5–0.8 pu sag while remaining connected, and (ii) after voltage recovers, restore active power to the pre-disturbance level within 1 second. Two characteristics of typical AC-DC PSU implementations make both parts difficult to meet today:

During the sag, standard PSU front-end protection (PFC stage) tends to trip within a few milliseconds rather than remaining connected and derating proportionally. From the grid's perspective, the load disconnects — the opposite of intended ride-through.

After voltage recovers, restarting a PSU from a tripped state is a cold-start sequence taking several seconds, which exceeds the 1-second restoration window.

These behaviours are not unique to any one manufacturer; they reflect typical AC-DC power supply design. Through the ERCOT NOGRR 282 process, the Data Center Coalition observed (Comments 282NOGRR-09 dated February 9, 2026):

"As the industry stands, some of the VRT regulations established in NOGRR282 are several years from being achievable ... the technology simply has not been developed yet."

ERCOT consequently adopted a phased applicability (Comments 282NOGRR-30 dated May 11, 2026, Sections 2.6.4(1) and 2.15(1)): a November 14, 2025 exemption reference date, an "if capable" basis through January 1, 2028, and strict compliance for facilities commissioned thereafter.

Recommendation: Drawing on the experience reflected in the ERCOT process, would IESO consider a similar phased path for §6.7.1(2) — facilities commissioned before a defined Ontario reference date applying the clause on an "if-capable" basis with documented fallback behaviour, and strict compliance for facilities commissioned thereafter?

IESO Response: From our discussions with developers, OEMs, utilities, and consultants, we have not identified any technical or implementation challenges related to the two aspects of the requirements mentioned.

Quote: "*During the sag, standard PSU front-end protection (PFC stage) tends to trip within a few milliseconds rather than remaining connected and derating proportionally. From the grid's perspective, the load disconnects — the opposite of the intended ride-through*."

There are no inherent equipment limitations requiring the rectifier or inverter to trip when the terminal voltage remains above 0.5 pu. The reduction of power consumption can be limited using a current limiter in the grid-side control system.

Quote: "*After voltage recovers, restarting a PSU from a tripped state is a cold-start sequence taking several seconds, which exceeds the 1-second restoration window*."

Our requirement specifies that the PSU shall remain connected and not trip. Instead, the rectifier or inverter may experience a momentary cessation and temporarily stop drawing power from the system when the terminal voltage falls below 0.5 pu. Once voltage recovers, this does not constitute a cold start. In fact, our customers have indicated that power restoration can occur almost instantaneously. The 1-second delay is intended as a confidence window to make sure that system voltage recovery is sustained before recovering the load.

That said, your concern is that these two technical features are not yet commercially available. Based on the experience reflected in the ERCOT process, mandatory compliance could be started after January 1, 2028. Given that few LCL projects are expected to enter service before that date, we do not see a need to establish a separate exemption period for these two technical aspects.

Issue 5 — §6.7.1: Cooling / Mechanical Load Treatment Under Deep Sag

§6.7.1 Table 1 requires uniform ride-through to 0.25 s between 0.2 and 0.5 Vnom, and 0.16 s below 0.2 Vnom. Cooling load is commonly 20–40% of total facility power and consists predominantly of induction motors. Induction motors below approximately 0.7 pu input voltage enter unstable regions of their torque-speed characteristic and risk stall, with consequential mechanical damage. This is a physical limit of the motor, not a control choice.

In the ERCOT NOGRR 282 stakeholder process, this issue surfaced through formal comments from the Data Center Coalition and others. The version recommended for approval by ERCOT's Technical Advisory Committee on April 29, 2026 (Comments 282NOGRR-30, Section 2.15(2)) included:

"Cooling or mechanical load at the LCEL facility may ride through or trip when voltage conditions are below 0.35 p.u. for any duration."

Section 2.15(3) further provided a recovery-baseline accommodation:

"For purposes of determining compliance with this requirement, if any cooling load at an LCEL facility were to trip for voltage conditions below 0.35 p.u. ... the amount of pre-disturbance cooling load would be subtracted from the total pre-disturbance consumption. This adjustment applies to the remaining requirements of this Section."

Recommendation: Would IESO consider an analogous approach — (a) permitting cooling and mechanical load to ride through or trip below approximately 0.35 pu, with the IT load remaining subject to the full §6.7.1 ride-through requirement; and (b) excluding tripped cooling load from the post-disturbance active power recovery target?

IESO Response: Quote: "*Induction motors below approximately 0.7 pu input voltage enter unstable regions of their torque-speed characteristic and risk stall, with consequential mechanical damage. This is a physical limit of the motor, not a control choice.*"

The statement captures an important physical effect of induction motors. However, the voltage threshold is not universally correct. We recognize that small motors serving cooling and mechanical loads may need to be tripped only when the POI voltage falls below 0.35 pu. Accordingly, we will add that due consideration will be given to the inherent limitations of such small motors, which may be permitted to trip under these conditions, provided that their trip does not restrict plant operation or delay the restoration of other IT loads following a system disturbance.

Please note that the fault ride-through requirement is not new in our system. We have previously experienced events where industrial loads tripped for out-of-zone faults, resulting in significant adverse impacts on system performance. In response, fault ride-through requirements were established in our system to mitigate such issues. Existing industrial loads are expected to meet these requirements.

Issue 6 — §6.7.1(2): "No Ramp" and Sub-Second Inrush Coordination

§6.7.1(2) requires restoration to pre-disturbance level within 1 second and states that no ramp rate limitations shall restrict recovery. At large data centre scale, thousands of distributed energy storage devices and PSU input bulk capacitors would recharge simultaneously upon voltage recovery if no coordination is permitted within the 1-second window. The aggregate inrush can stress PSU input rectification, busbars, and upstream switchgear, and may trigger protection — paradoxically delaying restoration. Individual soft-start limiters do not coordinate across the fleet; synchronous engagement aggregates to a large step at the POI.

A coordinated sub-second staggering within the 1-second window achieves full restoration on time while avoiding coincident inrush. This is not a ramp limitation on the recovery (which should remain prompt); it is coordination within the window. ERCOT NOGRR 282 Section 2.15(3)(a) similarly permits one control action per device per event.

Recommendation: Clarify that "no ramp rate limitations" prohibits ramp rates that would delay full restoration beyond the window and does not preclude sub-second coordinated staggering within it. Alternatively, the System Impact Assessment process could approve facility-specific recovery profiles that achieve full restoration within the window while avoiding POI step disturbance.

IESO Response: Thank you for pointing out this confusion. The load ramp rate limitation refers to the limit specified in Section 6.6, which applies under normal operation. This requirement does not preclude coordinated sub-second load staggering. We will clarify this in the next revision and note that ramp rate limitations necessary for equipment protection, such as measures to mitigate inrush current, may be implemented, provided that the required restoration time is achieved.

Issue 7 — §5 / §6.16 / Appendix A: Modelling Framework and Distributed Energy Storage Architectures

The IESO's connection assessment framework is model-driven: EMT, composite load models, and User Defined Models for PSS/E and DSA are required, with cross-platform benchmark consistency. The fidelity of these models depends on the underlying facility behaviour being deterministic at the time scales the model represents (sub-cycle to sub-second).

Where a facility's response is mediated by a small number of large, vendor-defined backup energy units, deterministic modelling is straightforward; the facility behaviour at the POI is reasonably predicted by the unit specifications. Where it is mediated by many distributed energy storage devices engaging independently based on local conditions (state of charge, thermal state, internal load, soft-start sequencing), deterministic modelling requires that those devices behave in a coordinated, specifiable way. Without coordination, per-device behaviour is stochastic, and the aggregate POI behaviour is difficult to predict in pre-commissioning study or validate in post-commissioning audit.

Recommendation: IESO consider whether its modelling requirements under §5, §6.16, and Appendix A should explicitly address facility designs in which behaviour at the POI is determined by a population of distributed energy storage devices — specifically, how the modelling framework expects coordination behaviour to be characterised so that the System Impact Assessment can rely on deterministic facility response.

IESO Response: At our meeting with you, we understood that distributed energy storage devices you are referring to are a modular UPS system composed of numerous small UPS units. As we clarified, the aggregation will not occur at the point of interconnection (POI). Typically, aggregation would take place at the low-voltage (LV) bus level, where similar UPS units connected to the same LV bus are aggregated as a single equivalent unit. This modelling approach is consistent with that used for battery energy storage systems (BESS) and renewable generation facilities.

In addition, we have been using deterministic models for system dynamic studies, even though real-time operations (e.g., renewable generation) are inherently stochastic. Deterministic models are designed to represent worst-case conditions across a range of operating scenarios. As such, we do not foresee any issues with continuing to apply deterministic modelling for these studies.

Independent Electricity System Operator
1600-120 Adelaide Street West
Toronto, Ontario M5H 1T1

Phone: 905.403.6900
Toll-free: 1.888.448.7777
E-mail: customer.relations@ieso.ca

ieso.ca

@OntarioIESO
linkedin.com/company/IESO

Connecting Today.
Powering Tomorrow.