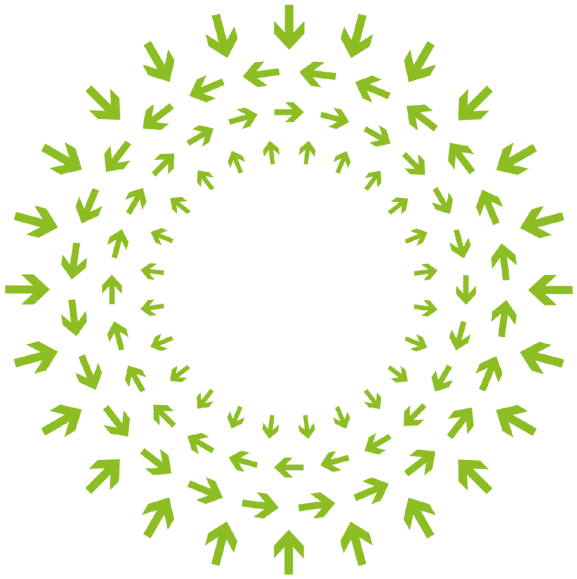


OPEN

Compute Project

Data Center Facility Project:

Interface Specification Between Hyperscalers and Utility Grid for Disturbance Performance

Version 1.0.0

OCP Base Specification
Effective September 17, 2026

Authors:

SN	Author	Affiliations
1	Mike Della Penna	Google
2	Nithish Mahalingam	Microsoft
3	Ashish Upreti	Google
4	Luz Belmont	Microsoft
5	Mauricio Cespedes	Meta
6	Sai Somayajula	NVIDIA
7	Slavko Vasilic	Meta

0. Version History

Date	Version #	Author	Description
Aug 31, 2026	1	See above	Final VRT Specification

Table of Contents

0. Version History 2
1. License 4
1.1. Open Web Foundation (OWF) CLA 4
2. Compliance with OCP Tenets 6
2.1. Openness 6
2.3. Efficiency 6
2.4. Impact 6
2.5. Scale 7
2.6. Sustainability 7
3. Introduction 9
4. Base Specifications 10
4.1. Description 10
4.2. Past and Upcoming Utility Requirements 10
4.2.1. Past Requirements 10
4.2.2. Upcoming Requirements 11
4.3. Detailed Requirements 12
4.4. Voltage Ride-Through and Post Fault Active Power Recovery 12
4.4.1. VRT & PFAPR - Requirement description 12
4.4.2. VRT & PFAPR - Overview of proposed utility specifications 13
4.4.3. Proposed specification 15
4.4.4. Additional Discussion 16
4.5. Implementation Plan Considerations and Outline 20
4.6. Future Work 21
4.6.1. Future requirements 21
5. References 22

1. License

THE UPDATED DEFAULT CONTRIBUTOR LICENSE AGREEMENT (CLA) IS **OWFa 0.9**. PLEASE VERIFY THE CORRECT CLA/FSA IS USED AND EXECUTED FOR THIS CONTRIBUTION.

1.1. Open Web Foundation (OWF) CLA

Contributions to this Specification are made under the terms and conditions set forth in **Modified Open Web Foundation Agreement 0.9 (OWFa 0.9)**. (As of October 16, 2024) ("Contribution License") by: **Google, Meta, Microsoft, NVIDIA**

Usage of this Specification is governed by the terms and conditions set forth in **Modified OWFa 0.9 Final Specification Agreement (FSA)** (As of October 16, 2024) **("Specification License").**

You can review the applicable Specification License(s) referenced above by the contributors to this Specification on the OCP website at https://www.opencompute.org/contributions/templates-agreements.

For actual executed copies of either agreement, please contact OCP directly.

Notes:

The above license does not apply to the Appendix or Appendices. The information in the Appendix or Appendices is for reference only and non-normative in nature.

NOTWITHSTANDING THE FOREGOING LICENSES, THIS SPECIFICATION IS PROVIDED BY OCP "AS IS" AND OCP EXPRESSLY DISCLAIMS ANY WARRANTIES (EXPRESS, IMPLIED, OR OTHERWISE), INCLUDING IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, FITNESS FOR A PARTICULAR PURPOSE, OR TITLE, RELATED TO THE SPECIFICATION. NOTICE IS HEREBY GIVEN, THAT OTHER RIGHTS NOT GRANTED AS SET FORTH ABOVE, INCLUDING WITHOUT LIMITATION, RIGHTS OF THIRD PARTIES WHO DID NOT EXECUTE THE ABOVE LICENSES, MAY BE IMPLICATED BY THE IMPLEMENTATION OF OR COMPLIANCE WITH THIS SPECIFICATION. OCP IS NOT RESPONSIBLE FOR IDENTIFYING RIGHTS FOR WHICH A LICENSE MAY BE REQUIRED IN ORDER TO IMPLEMENT THIS SPECIFICATION. THE ENTIRE RISK AS TO IMPLEMENTING OR OTHERWISE USING THE SPECIFICATION IS ASSUMED BY YOU. IN NO EVENT WILL OCP BE LIABLE TO YOU FOR ANY MONETARY DAMAGES WITH RESPECT TO ANY CLAIMS RELATED TO, OR ARISING OUT OF YOUR USE OF THIS SPECIFICATION, INCLUDING BUT NOT LIMITED TO ANY LIABILITY FOR LOST PROFITS OR ANY CONSEQUENTIAL, INCIDENTAL, INDIRECT, SPECIAL OR PUNITIVE DAMAGES OF ANY CHARACTER FROM ANY CAUSES OF ACTION OF ANY KIND WITH RESPECT TO THIS

SPECIFICATION, WHETHER BASED ON BREACH OF CONTRACT, TORT (INCLUDING NEGLIGENCE), OR OTHERWISE, AND EVEN IF OCP HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

2. Compliance with OCP Tenets

2.1. Openness

2.2. The Interface Specification working group exemplifies the Openness tenet through an open, transparent, and multi-company collaboration between Google, Microsoft, NVIDIA, and Meta. Developed in response to critical grid reliability challenges—such as the 1,500 MW simultaneous load loss event in August 2024—this specification establishes an open standard that bridges hyperscale data center operations with utility grid requirements. By publishing these interface guidelines and standardizing disturbance performance criteria, the working group provides a transparent framework that invites participation, feedback, and adoption across the broader Open Compute Project (OCP) community, utilities, and industry vendors, ensuring shared progress toward overall bulk electrical system stability.

2.3. Efficiency

Standardizing disturbance performance and interface requirements directly advances the OCP Efficiency tenet by streamlining design and deployment across the ecosystem. By establishing a unified baseline, hyperscalers and utility providers eliminate redundant, site-by-site engineering analysis and bespoke hardware adaptations. This shared alignment minimizes capital and operational overhead, accelerates interconnection timelines for massive AI clusters, and ensures consistent grid compliance without compromising operational continuity.

2.4. Impact

This document directly delivers on the OCP Impact tenet by driving transformative, industry-wide alignment between hyperscalers and utility grid operators:

- **Addressing Critical Industry Challenges:** By establishing a unified disturbance performance baseline, this specification directly solves the technical and operational hurdles of integrating AI/ML data centers (exceeding hundreds of megawatts) into modern utility grids.
- **Ensuring BES Grid Stability:** The defined interface requirements proactively prevent severe load loss events, safeguarding the stability, frequency, and overall reliability of the Bulk Electrical System (BES) during grid disturbances.
- **Standardizing Utility Requirements for Wider Adoption:** Standardizing explicit parameters for voltage ride-through and post-fault active power

recovery establishes an authoritative benchmark, driving broad adoption across hyperscalers, facility operators, and utilities globally.
- **Accelerating Technology Readiness:** Aligning design and operational criteria across the ecosystem unlocks rapid time-to-market and accelerates the deployment of next-generation AI/ML data center infrastructure.

2.5. Scale

By establishing a unified baseline for grid interconnection requirements, this specification directly serves as a catalyst for global industry scaling and supports the OCP Scale tenet in the following key ways:

- **Enabling Rapid Deployment of AI Clusters:** Standardizing disturbance performance criteria enables hyperscalers to pre-certify facility designs and rapidly deploy AI clusters (exceeding hundreds of megawatts) across diverse geographic regions without needing site-specific redesigns for varying regional grid standards.
- **Simplifying Internal Product Roadmaps:** Harmonizing conflicting regional grid requirements into a single OCP Base Specification simplifies and aligns internal product roadmaps for hyperscalers, data center operators, and equipment manufacturers worldwide.
- **Reducing Reliance on Bespoke Utility Hardware:** Standardizing clear operational expectations removes technical barriers and significantly reduces the reliance on custom, highly engineered utility-level hardware or bespoke facility retrofits to achieve grid compliance at scale.
- **Driving Broad Industry Adoption:** Providing a clear, multi-company reference standard encourages widespread adoption across hyperscalers, utilities, and vendors, establishing a repeatable blueprint that scales infrastructure seamlessly to meet expanding global power demands.

2.6. Sustainability

This document aligns with the Sustainability tenet in the following ways:

- **Mitigating Environmental Harm:** The primary objective of the specification is to maintain overall grid stability and reliability to prevent large load losses that can trigger widespread grid reliability issues and environmental harm. By effectively managing the dynamic demands of advanced computing infrastructure, it allows utilities to integrate large AI/ML data centers without requiring environmentally impactful, high-cost

utility-level hardware investments or bespoke, resource-intensive facility retrofits.

- **Efficient Power Management:** The document details explicit requirements for the Active Power Ramp Rate Limit and Post-Fault Active Power Recovery (PFAPR). These limits directly contribute to efficient power management by preventing sudden power swings and excessive operational strain on the grid's generation and transmission infrastructure.

3. Introduction

The growth of Artificial Intelligence (AI) and Machine Learning (ML) applications has led to the development of large-scale data centers, often exceeding hundreds of MW in power consumption. Coincidentally, an event in August 2024 in North America resulted in a simultaneous load loss of 1,500 MW from the Bulk Electrical System (BES) and more recently a similar event of even larger magnitude occurred in the Eastern Interconnection on July 22, 2026. A large load loss such as this resulted in the BES frequency and voltage rising enough to raise concerns for potential grid reliability issues. Utilities and regulatory bodies are also concerned about the impact these large loads may have on the BES. Over the last year, many entities including ESIG, NERC, and IEEE have highlighted a need for better understanding of the details of large load operations to inform load interconnection requirements that are under development. Utility and regulatory bodies claim that integrating these highly dynamic and power-intensive data centers into existing utility grids presents significant challenges and would require additional stringent interconnection requirements.

4. Base Specifications

4.1. Description

This **OCP Base Specification** establishes the interface requirements between hyperscale operators and the utility grid for performance during grid disturbances. The primary objective is to maintain grid stability and reliability through efficient power delivery that accommodates the dynamic demands of advanced computing infrastructure. The data center interface to the grid can happen at multiple voltage levels (most commonly medium voltage or high voltage), and this document assumes an agnostic view of connection voltage level and focuses on requirements that are applicable at data center load level at the point of interconnection (POI).

4.2. Past and Upcoming Utility Requirements

4.2.1. Past Requirements

Typical past large load utility requirements mostly focused on the following subcategories:

- **Capacity:** Transmission and generation capacity assessments to ensure load requested can be served reliably under a wide range of grid operational states
- **Protection and control:** Protection and control (P&C) design and coordination at the POI
- **Flicker and harmonics:** Power quality requirements for harmonics assessed at the POI under IEEE 1453 (flicker) and IEEE 519 (harmonics) standards
- **Power factor:** Specify a minimum power factor (e.g. 0.98 lagging) as measured at the POI
- **Load shedding:** Load shedding requirements (NERC PRC-006-2) were considered and deployed by some utilities for select sites

Due to organic and slower industry growth in the past, interconnection requirements were limited and many had flexibility in their exact application. With accelerating data center industry growth and the event of August 2024 noted above, data center integration has been a key focus for emerging utility requirements in the last year and a wide range of additional requirements have started to emerge.

4.2.2. Upcoming Requirements

At the highest level, utility and regulatory bodies are seeking to understand the possible impacts of emerging large loads on the bulk power system. Additional, and more stringent, requirements are being discussed and imposed. In addition, utilities are seeking additional detail about the expected data center load profiles and detailed models of the expected load behavior.

The Energy Systems Integration Group (ESIG) summarizes the areas of focus well in their report, *Large Load Performance Requirements: Current Practices and Recommendations* (ESIG, 2026), and for efficiency's sake, we repeat their summary here due to its high quality and comprehensiveness.

The ESIG Large Loads Task Force's investigation of large load characteristics and associated power system impacts identified reliability focus areas which informed the requirements analysis in this report[1]*:*

> ***Voltage ride-through:*** *a facility's ability to remain connected to the power system during and following a disturbance*
>
> ***Multiple disturbance ride-through:*** *large load facility designs that incorporate a disturbance counter, after which the large load would be switched to back-up supply or be disconnected*
>
> ***Active power recovery:*** *the post-fault recovery of load that was permitted to reduce during and immediately following a disturbance to facilitate voltage ride-through at a large load facility*
>
> ***Frequency ride-through:*** *a large load facility's ability to remain connected to the power system during high- and low-frequency excursions following a grid disturbance*
>
> ***Ramp rates, variability, and cycling:*** *large loads' potential to exhibit fast, large, and cyclical load changes, which can introduce frequency stability, voltage control, inter-area, and high frequency oscillation– related challenges*
>
> ***Reactive power:*** *the magnitude and control mode for large load*

[1] Energy Systems Integration Group. 2026. Large Load Performance Requirements: Current Practices and Recommendations. A report by the Large Loads Task Force. https://www.esig.energy/reports-briefs/large-load-interconnection-performance-requirements/

reactive power consumption

Voltage phase jump: *the ability of a facility to withstand sudden changes in phase following switching operations or disturbances*

Monitoring: *the required visibility into the operation of a large load facility considering data fidelity, resolution, and access*

Modeling: *the appropriate characterization of large load facilities based on the study type*

This specification focuses solely on disturbance performance requirements - primarily low voltage ride-through and post-fault active power recovery. It is expected that a follow-on specification will address active power fluctuation limits such as ramp rates and forced oscillations.

4.3. Detailed Requirements

The following sections feature a brief summary of the power system fundamentals of concern, details various specifications proposed by utilities and grid operators to date for that subcategory, briefly discusses the proposed specification for the requirement subcategory, and discusses mitigations possible to comply with the specification.

An Additional Requirements subcategory is also included to capture additional important requirements that don't fit neatly into the core list, but which nonetheless are critical for reliable utility operations.

For a more detailed discussion of the utility philosophy behind each of the requirements, please see ESIG's recent report, *Large Load Performance Requirements: Current Practices and Recommendations*.

4.4. Voltage Ride-Through and Post Fault Active Power Recovery

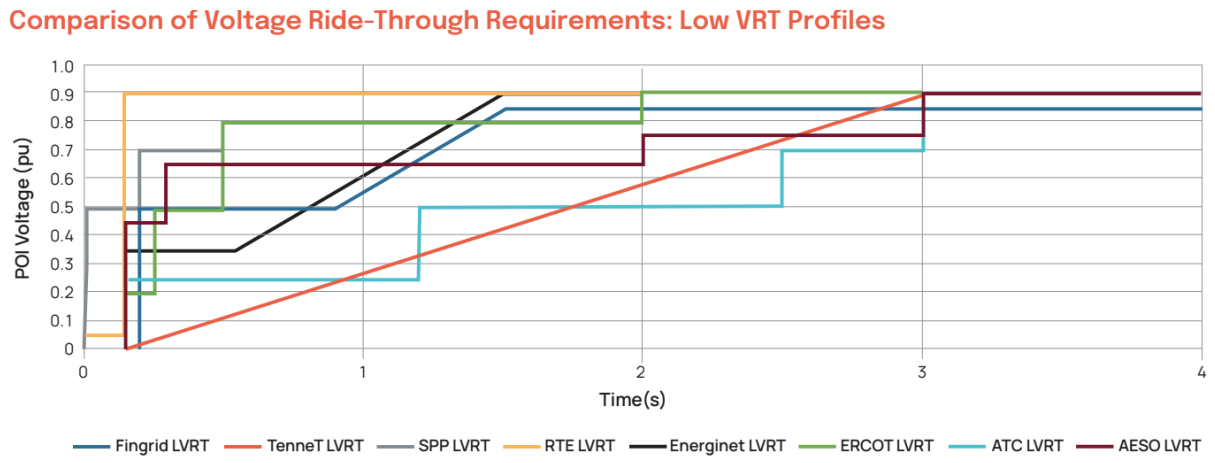
4.4.1. VRT & PFAPR - Requirement description

These requirements ensure that large data center loads remain connected to the grid during momentary voltage dips (or swells) rather than tripping offline. If massive blocks of demand disconnect unexpectedly during a grid fault, the resulting generation / load imbalance can cause frequency excursions and possibly lead to cascading grid

failures. As a component of the required ride-through capability, it is critical that large loads stay connected to the BES during all faults within the defined voltage-ride-through envelope and also don't disconnect based on a running count of faults over some period of time. Post-Fault Active Power Recovery (PFAPR) mandates a rapid, controlled return to pre-fault consumption levels once the voltage stabilizes above an acceptable threshold. This prevents the grid from experiencing a "secondary" disturbance caused by a sustained loss of load, ensuring the system returns to a balanced state within an acceptable period of time.

4.4.2. VRT & PFAPR - Overview of proposed utility specifications

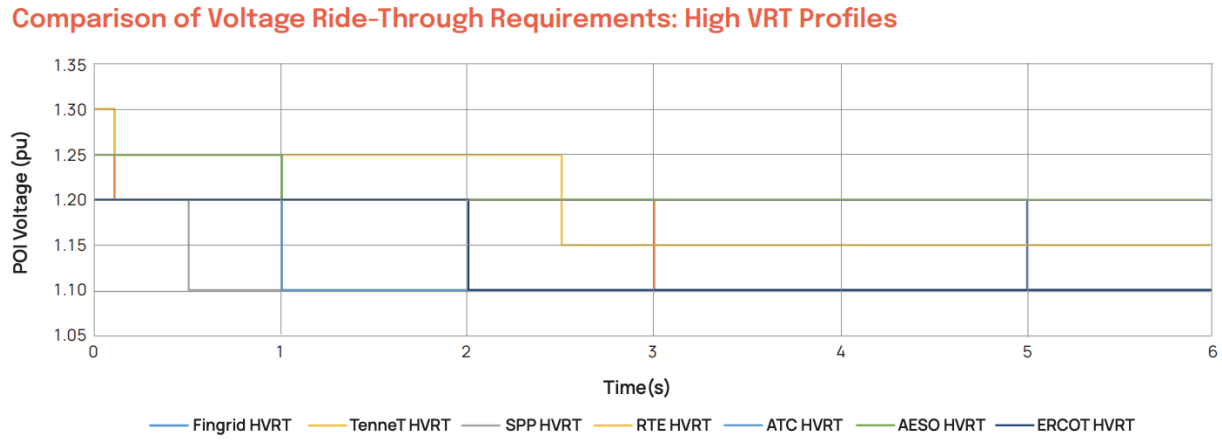
A number of utilities and grid operators have started to specify low/high voltage ride-through (LVRT/HVRT) curves as well as PFAPR requirements along with additional upcoming requirements. Industry status and proposed status has been documented in the IEEE white paper "Review of Industry Efforts and standards of grid readiness for data center deployment [1]". While an exhaustive discussion of each proposal is beyond the scope of this paper, a concise and graphical summary from the same ESIG report noted above is provided below.



Notes: AESO = Alberta Electric System Operator; ATC = American Transmission Company; ERCOT = Electric Reliability Council of Texas; LVRT = Low-voltage ride-through; RTE = Le Résea de Transport d'Electricité; SPP = Southwest Power Pool.

Source: Energy Systems Integration Group.

Figure 1. Proposed LVRT Requirements[2]



Notes: AESO = Alberta Electric System Operator; ATC = American Transmission Company; ERCOT = Electric Reliability Council of Texas; LVRT = Low-voltage ride-through; RTE = Le Résea de Transport d'Electricité; SPP = Southwest Power Pool.

Source: Energy Systems Integration Group.

Figure 2. Proposed HVRT Requirements[3]

[2] Figure reproduced with permission of ESIG
[3] Figure reproduced with permission of ESIG

4.4.3. Proposed specification

Below is a summary table outlining the HVRT, LVRT, and PFAPR performance requirements proposed as part of this specification. For this specification, the OCP hyperscalers shall adopt the convention proposed by MISO (MISO, July 2026) and commit to this LVRT/HVRT and PFAPR performance solely for the computational load component of the data center[4]**, which is typically ~70-90% of total data center load.** The requirements closely mirror ERCOT's proposal in NOGRR 282 (ERCOT, 2026) with some additional clarifications as noted.

Table 1. OCP Specification for Computational Load LVRT, HVRT, and PFAPR Performance

Line to neutral or line to line voltage as measured at POI **(p.u. of nominal voltage "V")**	**Minimum Ride-Through Time Before Transfer to Backup Allowed[5]** **(seconds)**	**Load Behavior During Ride-Through**	**Operational commitment once POI voltage returns to 0.90 ≤ V ≤ 1.10** **"PFAPR"**
V > 1.20	May transfer to backup instantaneously	N/A	N/A
1.10 < V ≤ 1.20	1.0	Maintain constant power	N/A
0.90 ≤ V ≤ 1.10	Normal range, continuous operation	Maintain constant power	N/A
0.80 ≤ V < 0.90	2.0	Continue consuming active power from the grid. May reduce power in proportion to V.	restore to 90% of pre-fault load within 2 seconds
0.50 ≤ V < 0.80	0.5	Continue consuming active power from the grid. May reduce power in proportion to V.	restore to 90% of pre-fault load within 2 seconds

[4] See section 4.4.4.4 for discussion on how how narrowing the focus to leaving out the cooling load component is not expected to generate second order effects on the computational load
[5] For voltage sags or swells of magnitude and duration outside these limits, computing loads are allowed to transfer to backup power schemes with no expectations on return timing/magnitude

0.35 ≤ V < 0.50	0.25	May reduce power to any value	restore to 90% of pre-fault load within 2 seconds
V < 0.35	0.15	May reduce power to any value	restore to 90% of pre-fault load within 2 seconds

4.4.4. Additional Discussion

4.4.4.1. Power consumption throughout HVRT/LVRT envelope

Consistent with proposals from ERCOT and MISO, the OCP hyperscalers commit to maintaining constant power consumption between 0.9-1.2 voltage p.u. Below 0.8 p.u. all the way to 0.5 p.u., it is allowable to reduce power in proportion to V. Below 0.5 p.u., power consumption is allowed to reduce to any level, *including zero MWs consumption*, as long as the PFAPR commitments noted can be met. It is worth noting that across all utility proposals seen by OCP hyperscalers, any level of power reduction is allowed when POI voltage falls below 0.5 p.u. (i.e. “100% current blocking”).

The last distinction is a critical point. Enabling true fault ride-through of the computational loads is not achievable with current OEM UPS and PSU architectures without risking additional operational and grid stability challenges. For example, truly riding-through and consuming at 0 p.u. requires larger capacitance to be added to the computational equipment to sustain load consumption during the fault window. Aside from hardware redesign impacts and often impractical retrofit requirements, adding this additional capacitance may also risk breaching the current consumption limit that grid operators such as ERCOT and EirGrid require not be exceeded. This creates a situation where solving for ride-through would introduce a new compliance risk upon recovery.

4.4.4.2. Voltage measurements

OCP hyperscalers propose that the voltage is measured at the POI and is either the line-to-neutral voltage or the line-to-line voltage. As certain types of faults may trigger line-to-neutral vs. line-to-line voltage voltage measurements that fall into different

buckets in Table 1 above, and these measurements are in part dictated by project-specific equipment specifications, OCP hyperscalers suggest the detailed implementation nuances of these schemes (for example, lesser of line-to-line and line-to-neutral, etc.) be addressed as part of the utility interconnection and project protection coordination efforts.

4.4.4.3. Reclosing schemes

The detailed implementation of utility reclosing schemes on the transmission network tends to vary slightly utility-by-utility. In certain cases, rapid reclosing over a short period of time post-fault can be interpreted by data center protective schemes as a deeper or longer fault. This challenge can be addressed in detailed project and utility specific coordination studies and protection system setting modifications, but such implementation across a wide fleet of geographically disparate sites is challenging and complex. More importantly, the lack of standardization of utility reclosing schemes introduces complexity that may result in less predictable data center responses to utility faults and subsequent recloser events. OCP hyperscalers suggest that additional utility efforts to standardize the specifics of transmission-level reclosing schemes over time will help deliver higher system reliability across the full system of transmission grid and data center load.

4.4.4.4. Cooling loads

It is important to note that mechanical/cooling loads will typically vary in their HVRT/LVRT/PFAPR behavior compared to computing loads. Most cooling loads for modern data centers have a variable frequency drive (VFD) interface and the LVRT/HVRT envelope differs meaningfully from that of rectifiers/power supply units (PSUs) and is most commonly set in compliance with IEEE 1566[6]. Post-tripping restoration times for cooling loads are longer than those for computing loads and can take minutes rather than seconds as some VFDs may require restart after interruption. Furthermore, any OEM modifications of the equipment requires a long lead time to make adjustments for the ride-through characteristics including updates to the existing standards.

[6] Typically allows shut down for voltages below 0.65 p.u. for over 500 ms. While IEEE 1566 applies to larger drives, smaller drives typically meet or exceed these requirements.

ERCOT has addressed this reality by modifying the calculation of the PFAPR pre-fault load to exclude cooling load amounts[7].

While firmware updates to enable enhanced LVRT/HVRT/PFAPR performance is likely a feasible path for computational load PSUs in many cases as noted below, OCP hyperscaler experience suggests cooling equipment firmware upgrades will not be sufficient to meet the ride-through requirements due to physics and hardware implementation. Generally speaking, enabling the required LVRT/HVRT/PFAPR performance for cooling loads is impractical to do without triggering hardware redesign or the addition of upstream energy storage, with the latter often not practical for existing sites with limited space. Adding to the challenge, while PSU OEMs and models are typically few in number, and in some cases, globally fungible, the OEM and equipment model landscape for cooling load equipment varies widely, adding significant execution complexity to any implementation plan.

Finally, hyperscaler cooling systems are already designed to sustain computational load operations throughout short-duration grid disturbances like open transitions. Even if cooling equipment momentarily trips and enters a restart sequence, computational loads typically can continue operating uninterrupted until facility cooling is restored given the presence and decay rate of thermal inertia in the facility. Accordingly, exempting cooling loads from the LVRT/HVRT/PFAPR fast-recovery requirements does not create a cascading reliability risk to computational load availability.

With this context, the OCP hyperscalers strongly support the cooling technology industry advancing fault ride-through capabilities of current and future VFD equipment designs to more closely align with the specifications noted above. However, given the early stage of these explorations and the expected complexities noted above, it is too soon to be able to commit to an implementation timeline for when the cooling load component can be compliant with LVRT/HVRT/PFAPR requirements as noted above.

[7] From NOGRR 282 language, “For purposes of determining compliance with this requirement, if any cooling load at an LECL facility were to trip for voltage conditions below 0.35 p.u. at the Service Delivery Point or POIB, the amount **of pre-disturbance cooling load would be subtracted from the total pre-disturbance consumption”**

4.4.4.5. Auxiliary loads

Auxiliary loads outside of computing and cooling loads are typically so small (< 1% for large scale data centers) that they can be safely ignored for system stability studies and this specification makes no clarifications or commitments with respect to the capabilities of this subset of total load.

4.4.4.6. Possible mitigations and implementation plan considerations

A range of mitigations are possible to comply with this specification, though the feasibility and implementation timeline associated with each option below varies significantly based on specific data center designs, equipment specifications, and project development/design timelines, which has important implications for industry when considering implementation plans and timing.

4.4.4.7. Rectifier / Power Supply Unit (PSU) Firmware Update

Data centers often deploy low-voltage (LV) rack-level rectifiers / PSUs to supply the IT equipment and the interface of this equipment characterizes the HVRT/LVRT/PFAPR capabilities of the computing load component. Firmware may be modifiable to allow stored energy in the PSU (typically by utilizing capacitors) to support the performance of the computational loads per the curves specified above.

This is an original equipment manufacturer (OEM) specific design question that needs to be investigated. Since this is broadly new territory and each PSU OEM has unique designs and capabilities, it is impossible to globally define feasibility and implementation timelines for such modifications. In addition, firmware updates are expected to vary by PSU models and hyperscalers may have a wide range of models used across the fleet with exact models selected based on specific computational workload and OEM supply availability. Utility implementation plans and grandfathering schemes should account for these realities when considering the go-forward application of new LVRT/HVRT/PFAPR requirements and additional discussion on this theme is included below.

4.4.4.8. MV-UPS Deployments

Certain data center architectures are exploring deploying in-line double conversion UPS schemes (MV-UPS) instead of the rack-mounted LV PSU design outlined above. In this case, it is also expected that firmware designs may allow stored energy in the MV-UPS to support the performance of the computational loads per the curves above. MV-UPS systems may also be considered as an addition to an existing facility to make a non-compliant facility compliant, but the feasibility of such an approach is heavily dependent on there being sufficient space available in the right spot(s) in the facility, practical electrical and controls interconnections that can be accessed without intrusive construction work, etc. In most cases, adding MV-UPS as a retrofit to make an existing facility compliant on a go-forward basis is judged to be too costly and disruptive to be a practical solution.

4.4.4.9. Grid-Side Mitigations

Utilities can help overall grid stability during disturbances by considering deploying dynamic reactive power support equipment in areas with high concentrations of data centers. This mitigation category doesn't help the data centers perform according to the performance criteria outlined above, but such additions could help reduce the severity of the observed voltage sag/swell at the facility POI, reducing the likelihood of an unstable outcome for the grid.
As an example, SVCs and E-STATCOMs can inject reactive power nearly instantaneously during and after a fault, reducing the depth of the voltage depression seen at the facility POI and accelerating post-fault voltage recovery. E-STATCOMs can be particularly effective in weak-grid conditions because they maintain full current injection capability even at severely depressed voltages.
Grid-scale SVCs and E-STATCOMs can be engineered to stabilize the broader transmission network by aggregating diverse, system-wide load profiles rather than addressing isolated, local demands.

4.5. Implementation Plan Considerations and Outline

Because enabling LVRT and PFAPR cannot be assumed to be a straight forward software or firmware update in all cases, the OCP

hyperscalers propose the following outline for a workable implementation plan:

- A specific effective date is established for the disturbance performance requirement
- Facilities that are already operational or that have executed interconnection agreements as of the effective date would not be required to comply for the lifetime of the facilities
- Growth at an already operational site above and beyond already executed interconnection agreement volume as of the effective date would be subject to the requirements, but *only for the incremental volume (i.e. incremental growth at an existing site doesn't trigger retroactive application of the requirements for previously grandfathered facilities)*
- Best-effort voluntary compliance pathways are established for grandfathered facilities and supported by utility engagement and possible incentives
- Facilities entering load study and contracting after the effective date would be expected to comply upon commencing operations

4.6. Future Work

This section outlines a handful of select additional requirements that are expected to be part of a future, more comprehensive, specification to be released in future versions of this specification as these are finalized between data center operators and Utilities.

4.6.1. Future requirements

- Disturbance performance
  - Frequency ride-through envelope
  - Rate of change of frequency
  - Phase jump
  - Counter-based protection schemes
- Active power fluctuations
  - Ramp rate
  - Intra-minute variation
  - Forced oscillations

5. References

5.1. IEEE Standards Association. (2026). *Review of industry efforts and standards of grid readiness for data center deployment* (IEEE White Paper). IEEE. https://standards.ieee.org/ieee/White_Paper/12415

5.2. MISO Energy, (2026), Interconnection Reliability Requirements for Large Loads, https://cdn.misoenergy.org/20260629%20LLWG%20Item%2007%20Draft_MISO_Large_Load_Interconnection_Requirement_Technical%20Compendium765704.pdf

5.3. ERCOT, (2026), NOGRR 282, https://www.ercot.com/mktrules/issues/NOGRR282#keydocs

5.4. Energy System Integration Group (ESIG), February 2026, Large Load Performance Requirements, https://www.esig.energy/reports-briefs/large-load-interconnection-performance-requirements/