

NERC

NORTH AMERICAN ELECTRIC
RELIABILITY CORPORATION

Characteristics and Risks of Emerging Large Loads

Large Loads Task Force White Paper

July 2025

RELIABILITY | RESILIENCE | SECURITY



3353 Peachtree Road NE
Suite 600, North Tower
Atlanta, GA 30326
404-446-2560 | www.nerc.com

Table of Contents

Preface	iv
Executive Summary.....	v
A New Challenge on the Horizon.....	v
NERC’S Response.....	v
Next Steps.....	vii
Chapter 1: Introduction to Large Loads	1
Intended Audience	1
NERC Large Loads Task Force	1
Large Load Definition	1
Background.....	2
Chapter 2: Characterization of Large Loads.....	3
Categories of Large Loads.....	3
Data Centers and Other Computational Load	3
Industrial Load	7
Hydrogen Production Facilities.....	8
Example Existing Large Load Programs and Constructs.....	9
Reliability Coordinators	9
Utility-Specific Programs.....	11
Future Large Load Definition Considerations.....	11
Single Number vs. Size Range	11
Behind-the-Meter Loads.....	12
Real-Time Load Behavior	13
Firm Load vs. Flexible Load	14
Chapter 3: Risks to the Bulk Power System	15
Large Load Observability and Data Risks.....	15
Long-Term Planning.....	15
Demand Forecasting.....	15
Resource Adequacy	16
Transmission Adequacy	17
Operations/Balancing.....	18
Short-Term Demand Forecasting.....	18
Balancing and Reserves	19
Lack of Real-Time Coordination.....	20


Outage Coordination in Operations Planning.....	21
Stability.....	21
Ride-Through	22
Frequency Stability	23
Rotor Angle Stability	25
Voltage Stability	27
Resonance- and Converter-Driven Stability.....	28
Forced Oscillations.....	29
Power Quality	32
Harmonics	32
Voltage Fluctuations	34
Physical and Cyber Security Risks	36
Cyber Security	36
Load-Shedding Programs and System Restoration	36
Manual Load-Shed Obligation Impacts.....	36
Automatic Under-Frequency Load-Shedding Programs	36
System Restoration.....	36
Chapter 4: Conclusion	38
Appendix A: Large Load Construct Data	40
Appendix B: Comparison of Characteristics and Risks	41
Appendix C: Acknowledgements	42

Preface

Electricity is a key component of the fabric of modern society and the Electric Reliability Organization (ERO) Enterprise serves to strengthen that fabric. The vision for the ERO Enterprise, which is comprised of NERC and the six Regional Entities, is a highly reliable, resilient, and secure North American bulk power system (BPS). Our mission is to assure the effective and efficient reduction of risks to the reliability and security of the grid.

Reliability | Resilience | Security
Because nearly 400 million citizens in North America are counting on us

The North American BPS is made up of six Regional Entities as shown on the map and in the corresponding table below. The multicolored area denotes overlap as some load-serving entities participate in one Regional Entity while associated Transmission Owners/Operators participate in another.



MRO	Midwest Reliability Organization
NPCC	Northeast Power Coordinating Council
RF	ReliabilityFirst
SERC	SERC Reliability Corporation
Texas RE	Texas Reliability Entity
WECC	WECC

Executive Summary

A New Challenge on the Horizon

As the North American power system evolves, loads such as computational, industrial, and hydrogen production facilities are seeking to connect to the bulk power system (BPS) faster than ever before and at a magnitude beyond the largest currently operating loads. There is already evidence that large loads impact BPS reliability. For example, the Eastern and Electric Reliability Council of Texas (ERCOT) Interconnections have observed load-reduction events¹ with each Interconnection experiencing approximately 1,500 MW of voltage-sensitive load reduction. The event in the Eastern Interconnection was primarily attributed to data centers and other power electronic loads (PEL) transferring load to backup generation and caused frequency overshoot and high voltages. The ERCOT Interconnection event involved many different types of loads of varying size reducing consumption during an extended low-voltage period in West Texas due to a protection system misoperation. These load-reduction events highlight some of the potential risks posed by large loads utilizing the BPS and why NERC is closely examining this issue.

There is already evidence that large loads impact bulk power system reliability.

NERC'S Response

NERC's Reliability and Security Technical Committee (RSTC) established the Large Load Task Force (LLTF) to analyze the reliability impacts related to emerging large loads (e.g., data centers (including cryptocurrency mining, and artificial intelligence (AI)), industrial facilities, and hydrogen production facilities). This white paper characterizes large


loads and defines the reliability risks that they may pose to the BPS. A reliability gap analysis white paper will follow, shedding more light on where the most significant risks exist.

To better understand the reliability impacts, NERC established the Large Loads Task Force.

The loads examined in this paper range in size from several megawatts to several gigawatts, posing novel challenges to the reliability and security of the BPS. Several Reliability Coordinators (RC) and utilities have existing large-load constructs based primarily on facility peak demand. However, additional characteristics should be considered in any definition of large loads, as this white paper shows that peak demand is only one of many characteristics that can impact BPS reliability.

¹ "Incident Review - Considering Simultaneous Voltage-Sensitive Load Reductions," NERC, Jan. 2025. Available: https://www.nerc.com/pa/rrm/ea/Documents/Incident_Review_Large_Load_Loss.pdf

The LLTF analyzed the following risk categories for this white paper:




The following recommendations are offered as guidance for future work to ensure the reliability and security of the BPS:

- ✓ **Recommendation 1:** The NERC LLTF should identify existing processes and standards that do not fully address the risks of emerging large loads, as planned for the LLTF’s second work item—*White Paper: Assessment of Gaps in Existing Practices, Requirements, and Reliability Standards for Emerging Large Loads*.
- ✓ **Recommendation 2:** The LLTF should identify potential mitigations to risks posed by emerging large loads through improvements to existing planning and operation processes and interconnection procedures for large loads as planned for the LLTF’s third work item—*Reliability Guideline: Risk Mitigation for Emerging Large Loads*.
- ✓ **Recommendation 3:** The LLTF should clearly define each of the identified characteristics of emerging large loads and develop a framework for classifying large loads.
- ✓ **Recommendation 4:** The NERC Load Modeling Working Group should create and approve load models that can show the characteristics and risks of each category of emerging large loads in simulations.
- ✓ **Recommendation 5:** The NERC System Protection and Control Working Group should assess possible protection system impacts to the BPS from emerging large loads.
- ✓ **Recommendation 6:** The NERC Energy Reliability Assessment Working Group and Probabilistic Assessments Working Group should investigate methods for grid operators and planners to assess the risks potentially posed by emerging large loads to resource adequacy.

Next Steps

To address the near- and long-term risks associated with large loads, NERC is identifying gaps in integrating these large loads onto the grid, enabling a coordinated approach to address this issue effectively. This white paper is expected to be published in Q4 2025.



Chapter 1: Introduction to Large Loads

Intended Audience

This white paper is intended for the following NERC registered entities, external entities, and broader groups:

- Planning Coordinators (PC)
- Transmission Planners (TP)
- Transmission Owners (TO)
- Transmission Operators (TOP)
- Distribution Providers (DP)
- Balancing Authorities (BA)
- Reliability Coordinators (RC)
- Large load developers, owners, operators, or other related companies
- Generator Owners (GO)
- Generator Operators (GOP)
- Reliability and Security Technical Committee (RSTC) subgroups

This white paper identifies, validates, and prioritizes the characteristics and risks of emerging large loads to the BPS. Entities responsible for operating and planning facilities on the BPS should be aware of these characteristics and risks to maintain reliable service. Subsequent LLTF work items, like the next white paper and reliability guideline, will need to address the risks noted in this white paper. Finally, this white paper identifies areas where potential security risks associated with emerging large loads require further assessment by related RSTC subgroups.

NERC Large Loads Task Force

The purpose of the NERC LLTF is to better understand the reliability impacts that emerging large loads, such as data centers and other computational loads, large industrial loads, and hydrogen production facilities, will have on the BPS. The LLTF has two primary phases identified in the NERC LLTF scope² as follows:

- Phase 1: Identify unique characteristics and risks of large loads
- Phase 2: Identify gaps and potential risk mitigation

Large Load Definition

NERC's LLTF was tasked with defining large loads. The definition was created for use in this white paper and the other work products from NERC's LLTF. This definition is expected to be modified in the future for specific applications, like the *NERC Glossary of Terms Used in NERC Reliability Standards*.³ For the purposes of this white paper and the following white paper and reliability guideline, large loads shall be defined as the following:

"Any commercial or industrial individual load facility or aggregation of load facilities at a single site behind one or more point(s) of interconnection that can pose reliability risks to the BPS due to its demand, operational characteristics, or other factors. Examples include, but are not limited to, data centers, cryptocurrency mining facilities, hydrogen electrolyzers, manufacturing facilities, and arc furnaces."

² "Large Loads Task Force (LLTF)," NERC, Aug. 2024. [Online]. Available: <https://www.nerc.com/comm/RSTC/LLTF/LLTF%20Scope.pdf>

³ "Glossary of Terms Used in NERC Reliability Standards," NERC, Feb. 2025. Available: https://www.nerc.com/pa/Stand/Glossary%20of%20Terms/Glossary_of_Terms.pdf

The LLTF has not yet provided a megawatt (MW) threshold or more granular characteristics of large load as further analysis and development is forthcoming.

Background

Emerging large loads have grown substantially in recent history, and forecasts predict that data centers alone may account for as much as 12% of all U.S. electricity consumption by 2028⁴ (up from 4.4% in 2023). In addition to the rapid, demonstrated, and forecasted growth of emerging large loads, the characteristics that inform their electrical behavior are not as predictable as conforming loads. The electrical behavior of emerging large loads during both normal and abnormal grid conditions represents new challenges for planners and operators. Operating experience is limited with emerging large loads. Large loads may come in different forms, although the rapid growth associated with data centers have led the authors of this white paper to focus on the characteristics and potential risks associated with integrating these types of users of the BPS. This white paper will discuss emerging large loads' characteristics and how each one informs specific risks to the BPS.

⁴ A. Shehabi *et al.*, "2024 United States Data Center Energy Usage Report," Lawrence Berkeley National Laboratory, Dec. 2024. Available: <https://eta-publications.lbl.gov/sites/default/files/2024-12/lbnl-2024-united-states-data-center-energy-usage-report.pdf>

Chapter 2: Characterization of Large Loads

This chapter categorizes the large loads that have emerged in recent years across North America, reviews existing large load programs and constructs in the United States and Canada and addresses certain factors that could inform a formal NERC large load definition.

Large load customers are connecting to the grid faster than ever before. These loads range widely in size: some are facilities of several MWs that can connect at the distribution level, while others are multi-gigawatt (GW) complexes that must be served by the transmission system and may require new transmission facilities or other upgrades to support the increased system demand.

In addition to their size, large loads have other characteristics that pose reliability risks to the BPS. For example, many new large loads are PELs that use software controls to manage the load behavior rather than the mechanical controls used by older motor loads. PEL control systems enable loads to rapidly change electricity consumption, creating previously unseen load behavior that can cause BPS planning and operational challenges.

Categories of Large Loads

The following large load categories are based on the function or activity of the facility in question. Each type will require tailored considerations for interconnection, BPS reliability, and grid integration.

Data Centers and Other Computational Load

A data center is a physical room, building, or facility that houses IT infrastructure for building, running, and delivering applications and services.⁵ Data centers are among the fastest-growing energy consumers in North America, supporting cloud computing, AI, cryptocurrencies, and other digital services. They range from hyperscale facilities operated by tech giants, like Meta's planned 2 GW AI training facility in Louisiana,⁶ to smaller data centers that serve multiple enterprises. Some data center loads are known to have pulsed and non-linear characteristics, giving them extremely quick ramping potential that can introduce power quality and stability issues. Cryptocurrency mines are generally not classified as data centers though much of their makeup is similar. Cryptocurrency mines may have different types of IT equipment compared to traditional and AI data centers, and they may not prioritize uptime as much as a traditional data center.

Data centers and other computational loads are PELs that consist of high-performance computing (HPC) systems in which large clusters of computing resources work in concert to perform complex analysis and other tasks. These loads are characterized by high energy consumption, variable operational demand, and significant cooling requirements. "Data center" is a term that encompasses many different end uses and configurations, as explained in this section's various sub-sections. The internal configuration of the load can be very complex, including internal protection and backup power systems. Data centers and other computational loads are often located in areas with access to reliable and cost-effective electricity, requiring carefully tailored grid planning and interconnection strategies.

⁵ S. Susnjara and I. Smalley, "What is a data center?," IBM Think, Sep. 04, 2024. <https://www.ibm.com/think/topics/data-centers> (accessed May 27, 2025).

⁶ "Meta's Richland Parish Data Center." Available: <https://datacenters.atmeta.com/wp-content/uploads/2024/12/Metas-Richland-Parish-Data-Center.pdf>

The four main components in data centers are listed as follows:

- **IT-related equipment** consists of servers, storage, and networking gear driven by power electronics. These components generally comprise at least half of the data center's energy consumption.^{7,8} Data center industry experts state that IT-related equipment may make up 60–95% of a data center's total energy demand.
- **Power delivery systems** control and distribute power in the facility. They may consist of multiple sources of external power, internal power distribution systems, emergency backup systems (including uninterruptible power supply (UPS) and on-site generation), and sophisticated control systems that ensure continuity of service to their critical network equipment. Data center industry experts note that power delivery systems for HPC loads (AI data centers) commonly exclude UPS protection systems for the IT equipment. Instead, they employ a checkpoint process where checkpoints are saved at regular intervals. If there is an interruption to the IT equipment's power, the checkpoint may be restored after the power interruption ends.
- **Cooling systems** typically consist of PEL including variable speed drives, inverters, and other internal equipment that can create harmonics and increased reactive power demands.
- **Miscellaneous** lighting and security loads are necessary for operation of the load. This typically makes up the smallest portion of the total energy usage of the load.

Figure 2.1 shows how these components and other features, such as UPS and backup generation, may be arranged in a data center. Notably, the desired uptime of the facility determines what tier of backup equipment is required. This figure is more relevant to traditional data centers. For AI training loads, there may be no UPS for the IT equipment.




Figure 2.1: Tier I–IV Data Center One-Line⁹

⁷ J. Davis, "Large data centers are mostly more efficient, analysis confirms," *Uptime Institute Blog*, Feb. 07, 2024.

<https://journal.uptimeinstitute.com/large-data-centers-are-mostly-more-efficient-analysis-confirms/>

⁸ "Technical Update on Load Modeling. EPRI, Palo Alto, CA: 2023. 3002027016"

⁹ S. Chalise et al., "Data center energy systems: Current technology and future direction," 2015 IEEE Power & Energy Society General Meeting, Denver, CO, 2015, pp. 1-5, doi: 10.1109/PESGM.2015.7286420.

Traditional Data Centers

Traditional data centers power the digital services of the modern world. They enable almost every online service, including banking, healthcare, e-commerce, and social media.¹⁰ Traditional data centers make up most of the currently operational data centers and exclude AI training and inference facilities. The exact configuration and business model can vary greatly, but most traditional data centers have similar electrical characteristics. Traditional data centers have historically been characterized by their smaller size¹¹ (under 30 MW in many cases) and limited variability. Traditional data centers commonly prioritize up time – the percentage of time that a data center provides uninterrupted service. These types of data centers are known for their redundancy. There are a few organizations, like the Uptime Institute,¹² that classify data centers into Tiers I-IV. Tier IV data centers have the highest redundancy, resulting in uptimes of over 99.995%.

Some companies, like utilities, use small local data centers to meet their business needs. These types of data centers are typically very small (less than 1 MW) and are not the focus of this white paper.

Larger businesses may use self-hosted data centers for their own business use cases. Social media is a prime example of this, as it takes many large data centers to host services that are used by millions of people. These data centers are owned and operated by a single company, so they will have the most information available about physical infrastructure, electrical characteristics, and demand patterns.

The last major type of traditional data center is the cloud data center, which may rent out its computing power to multiple businesses on a recurring or temporary basis. These data centers have information about the physical equipment and characteristics but may not have as much information on demand patterns as a self-hosted, self-use data center.

Artificial Intelligence Training Data Centers

Data centers powering AI technologies are among the newest types of load customers. AI is generally defined as technology that enables computers to perform complex tasks like translating spoken and written language, analyzing data, and making recommendations.¹³ AI training refers to the creation and modification of an AI model and its associated model weights.

AI facilities (training and inference) include powerful, application-specific integrated circuits (ASIC) or graphics processing units (GPU). Current industry trends, especially in AI applications, show GPUs getting smaller with higher power ratings than historic GPUs. These computing components operate best in very specific conditions, requiring specific temperatures and humidity. This is often a key consideration when siting data centers. If temperature and humidity ranges are too vast, the temperature and humidity control equipment may be too burdensome. Cooling load demand can vary widely depending on the specific data center's cooling equipment and based on the intensity of the processing tasks that the IT equipment is assigned at any given time. The cooling load requirements are highly dependent on the end use. For instance, cooling load has different characteristics for AI data centers compared to traditional data centers.

The electrical demand is different for training an AI model versus using an AI model to create inferences.¹⁴ Generally, AI training executed in large clusters has rapid fluctuations during training periods and while saving checkpoint

¹⁰ "What is a datacenter? - Microsoft Datacenters," Microsoft Datacenters, May 05, 2025.

<https://datacenters.microsoft.com/whatisadatacenter/>

¹¹ B. Srivathsan, M. Sorel, and P. Sachdeva, "AI power: Expanding data center capacity to meet growing demand," *McKinsey & Company*, Oct. 29, 2024. <https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/ai-power-expanding-data-center-capacity-to-meet-growing-demand>

¹² "Tier Classification System," Uptime Institute. <https://uptimeinstitute.com/tiers>

¹³ Google Cloud, "What Is Artificial Intelligence (AI)?," *Google Cloud*, 2025. <https://cloud.google.com/learn/what-is-artificial-intelligence>

¹⁴ Trends in AI inference energy consumption: Beyond the performance-vs-parameter laws of deep learning, <https://doi.org/10.1016/j.suscom.2023.100857>

progress. Additionally, the transition from training to saving checkpoint progress (or vice versa) may happen in under one second. See [Figure 2.2](#) and [Figure 2.3](#), which shows the demand from a 50 MW block of a larger 200 MW AI training data center that starts a training run. The data was recorded with a high-speed recorder and was supplied by EdgeTunePower.¹⁵ In the fastest ramping period, the demand changes at a rate of 1.9 p.u. per second for about 250 milliseconds.

While the training is active, jittery spikes up and down occur, shown between 360 and 660 seconds in [Figure 2.2](#). Google explains the problem in a Google Cloud blog post,¹⁶ stating that large power spikes are observed when transitioning between “idle and peak utilization levels” but that they also occur “when the workload is running normally, mostly attributable to alternating compute- and networking-intensive phases of the workload within a training step.”




Figure 2.2: An AI Training Data Center Begins a Training Run (EdgeTunePower)

¹⁵ “About Us | EdgeTunePower,” EdgeTunePower. <https://www.etpower.ca/about-us>

¹⁶ H. Gan and P. Ranganathan, “Balance of power: A full-stack approach to power and thermal fluctuations in ML infrastructure,” *Google Cloud*, Feb. 11, 2025. <https://cloud.google.com/blog/topics/systems/mitigating-power-and-thermal-fluctuations-in-ml-infrastructure> (accessed May 29, 2025).




Figure 2.3: Zoomed-In Views of AI Training Data Center Ramping Up and Down Due to Transitioning Between Training and Saving Checkpoint Progress

Artificial Intelligence Inference Data Centers

AI inference is the process of using a pre-trained AI model to generate output(s) based on new input. AI inference uses similar, high-power GPUs when compared to AI training. It is common for data centers to be purpose-built for either inference or training workloads. The primary difference is in the underlying computations happening on the GPUs. The rapid ramps up and down seen in AI training/checkpoint cycles are not observed in current AI inference methods. Some companies in the AI industry, like NVIDIA, postulate that inference will be the greater electrical demand in the future when compared to training. In February 2022, Google stated that 60% of its total energy use went toward inference and 40% went toward training.¹⁷ However, this was before some of the largest growth in large language models (LLM) and AI over the past three years.

Cryptocurrency Mining

Cryptocurrencies (e.g., Bitcoin) rely on “data mining,” wherein computers solve mathematical puzzles to validate and enable various transactions. Cryptocurrency mining facilities consist of purpose-built servers and mining units—typically ASIC miners. Depending on the size of the operation and the number of units, cryptocurrency mining can consume large amounts of electricity to power the machines and cool the equipment.

Unlike data centers, which ramp up or down in response to customer demands or training cycles, the power consumption of cryptocurrency mining facilities is more stable and typically driven by internal computational needs, giving cryptocurrency miners flexibility in responding to system conditions, like reducing demand when prices are high.

Industrial Load

Industrial loads consist of facilities and equipment used for producing, processing, or assembling goods. These operations often require robust electric infrastructure and are typically more complex and energy-demanding than residential or commercial loads. Industrial loads are commonly referred to as non-conforming loads. This sector is set

¹⁷ D. Patterson, “Good News About the Carbon Footprint of Machine Learning Training,” *Google Research*, Feb. 15, 2022. <https://research.google/blog/good-news-about-the-carbon-footprint-of-machine-learning-training/> (accessed May 28, 2025).

to grow as more industrial activity moves to North America and existing facilities replace oil and gas with electricity as a primary fuel source.

Mining and Mineral Processing

The mining and mineral processing industries extract and refine raw materials essential for sectors including hard rock mining, coal mining, mineral processing, and cement production.

Metals and Heavy Manufacturing

Metals and heavy manufacturing consists of industries involved in the processing of metals and large-scale manufacturing, including steel mills, smelters and metal refineries, rolling mills, foundries, and casting facilities.

These industries use electricity-intensive processes, including electric arc furnaces (EAF), induction furnaces, electro-refining, arc welding and plasma cutting, and large-scale electrolysis.

Semiconductor and Electronics Manufacturing

These industries require precision power and cleanroom environments. Examples include semiconductor fabrication, liquid-crystal display (LCD) and organic light-emitting diode (OLED) and display manufacturing, battery manufacturing, and printed circuit board (PCB) production.

Chemical and Petrochemical Processing

These industries rely on high-energy chemical reactions to facilitate their processes. Examples include ammonia and fertilizer plants, chlor-alkali and electrochemical processing, refining and petrochemicals (oil refineries, synthetic fuels, plastics), and carbon capture and utilization (CCU).

Oil and Gas Production

The production of oil and gas requires vast amounts of energy. To improve oil and gas field equipment efficiencies and reduce costs, many operators are switching equipment to run on electricity and not fossil fuels. The buildout of liquefied natural gas plants is also driving load growth in certain regions of North America.

Hydrogen Production Facilities

This white paper focuses solely on hydrogen production facilities that use electricity to perform electrolysis. Hydrogen is an emerging option in the push to decarbonize the grid. The hydrogen production process begins at a hydrogen production plant, where water is treated and purified to remove dissolved solids, organic compounds, and other contaminants, resulting in pure water (H₂O).

This purified water then undergoes electrolysis, which splits the water into hydrogen and oxygen molecules. Each electrolyzer used in this process typically consumes 5–10 MW of power, and multiple units can operate within the same facility, as shown in [Figure 2.4](#). After electrolysis, the hydrogen is cooled and liquefied. Liquid hydrogen can then be stored or transported for use in various applications, including fueling vehicles (e.g., cars, trucks, airplanes), powering data centers, providing backup power, and supporting manufacturing processes.

Hydrogen electrolyzer facilities under development could reach the multi-gigawatt scale. Initial surveys have shown that over 85% of the facility load consists of power electronic converters with the remaining load being motor driven for H₂ compression, water treatment and cooling, and other plant equipment.




Figure 2.4: Example Hydrogen Production Facility One-Line Diagram (EPRI)

Example Existing Large Load Programs and Constructs

Several entities in North America have requirements for large load interconnection and/or preliminary definitions and programs for large loads. Notably, the only organizations included in this white paper are ones that voluntarily provided information to the writers of the white paper. Of the RCs and utilities with existing large load constructs, many of the current constructs are based on transmission facility ratings with no additional accounting for the potential impact that large loads may have on the BPS. Utility-specific programs not detailed in this section are described in [Appendix A: Large Load Construct Data](#). All information about programs and constructs in this white paper reflects current practices and is subject to change.

Reliability Coordinators

ERCOT

In the ERCOT Interconnection, a large load is defined as 75 MW or larger. At this size, transmission service providers (TSP) in ERCOT typically must develop and construct transmission upgrades to support the full amount of load that is seeking interconnection.

ERCOT established an interim large load interconnection process through Market Notice W-A032522-01¹⁸ to ensure compliance with existing NERC requirements and for the reliable interconnection of large loads to the ERCOT system. TSPs are required to submit interconnection studies that meet the requirements of NERC Reliability Standard FAC-002-2 for each applicable large load proposing to interconnect to the ERCOT system.

¹⁸ "Market Notice - Interim Large Load Interconnection Process," *Ercot.com*, Mar. 25, 2022.
https://www.ercot.com/services/comm/mkt_notices/W-A032522-01

ERCOT's interim large load program applies to projects seeking to interconnect in two years or less that meet one of the following criteria:

- For standalone loads not co-located with a generation resource:
 - A new load is requesting a total peak demand of 75 MW or more
 - An existing load is requesting an increase of its total peak demand by 75 MW or more
- For loads co-located with a generation resource:
 - A new load is requesting a total peak demand of 20 MW or more
 - An existing load is requesting an increase of its total peak demand by 20 MW or more

NYISO

To supplement the load interconnection procedures in the New York Independent System Operator (NYISO) Open Access Transmission Tariff (OATT), NYISO has defined load interconnection procedures in its *Transmission Expansion and Interconnection Manual*. These procedures specify which load interconnections are subject to NYISO's interconnection studies: those with a load of 10 MW or greater connecting at a voltage level of 115 kV or above or a load 80 MW or greater connecting at a voltage level below 115 kV. Such projects must be evaluated to determine whether they may degrade system reliability or adversely affect the operation of the New York state transmission system.

NYISO performs a single interconnection study, the System Impact Study, for load interconnections, the scope and results of which are reviewed by NYISO's operating committee. If the System Impact Study indicates the potential for an adverse reliability impact, NYISO identifies the need for potential upgrades required to reliably interconnect the load project. Should the load customer elect to proceed, required upgrades are evaluated by the applicable TO(s), and the engineering, procurement, and construction details as well as the cost responsibility for such upgrades are memorialized in Interconnection Agreements between the load customer and the applicable TO(s).

Projects that do not meet the above MW/voltage thresholds are subject entirely to the connecting TO(s) load interconnection procedures; NYISO is not involved in any interconnection studies for such projects.

Southwest Power Pool (SPP)

The Delivery Point Addition Process in SPP, defined under Attachment AQ of the SPP tariff, governs how new or modified delivery points are added to the transmission system. A delivery point is where power is transferred from the grid to a local utility or load-serving entity. When a new delivery point is proposed or an existing one is changed, the transmission customer (TC) submits a delivery point assessment (DPA) request. The transmission provider, SPP, then performs an engineering study to assess reliability impacts, and the host TO performs a load connection study (LCS) to assess load interconnection needs. SPP reviews the LCS, coordinates with stakeholders, and determines if system upgrades are needed. Costs for network upgrades are regionally funded, and interconnection facility upgrade costs are agreed upon between the TC and TO. Once approved, the delivery point and associated upgrades are added to the TC's service agreement, ensuring that the agreement is recognized in planning, operations, and billing. This process ensures reliability, coordination, transparency, and fair cost allocation across the SPP system.

ISO-NE

The Independent System Operator New England (ISO-NE) does not have an explicit "large loads program." However, ISO-NE learns of new loads seeking to interconnect through either a transmission service application for new load (regional network load or local network load) under regional network service (typically 115 kV or above) or local network service (typically below 115 kV, not including distribution facilities), depending on how the new load would be served. However, large data center loads have not yet sought interconnection in the area through application for these services, and there is no definition of "large loads."

Only applications for wholesale service are submitted to ISO-NE. Applications for retail service to serve load are submitted to the New England TOs, and ISO-NE has no visibility into those applications.

Utility-Specific Programs

In this category, Dominion Energy has the most experience with large load interconnections because Virginia (which is within Dominion's service area) is home to the world's largest concentration of data centers. Other known utility-specific programs are described in [Appendix A: Large Load Construct Data](#).

Dominion

Per Section 2.13 End User Facilities (Load Interconnection) of the Electric Transmission Facility Interconnection Requirements End Users Facilities,¹⁹ Dominion sets a 100 MW threshold for loads tapping the transmission system at a single point of interconnection. Below 100 MW, Dominion allows the load to be tapped with only isolation switches; above 100 MW, a ring bus configuration is required. As part of Dominion's planning criteria, the amount of load connected to a single substation is limited to 300 MW and requires no more than 300 MW load loss for N-1 and N-1-1 contingencies.

For any load connection to the transmission system, Dominion requires customers to submit a detailed request form with specific sections for data centers that require information about ride-through capabilities and other features. The Dominion 500 kV system is reserved for bulk power transfer, and load interconnections are not allowed to reach this voltage. Most data center load connections are on the 230 kV system, but some occur on the 115 kV and 138 kV systems.

Future Large Load Definition Considerations

Given the numerous categories and subtypes of large loads, a classification system based solely on a load's peak demand may be overly simplistic. Due to the differences highlighted above, it cannot be assumed that all loads will behave similarly. Several other factors, such as load ramp rates, real-time behavior and flexibility, protection systems, and backup power schemes, should be considered if a large load definition is developed for regulatory procedures or policies. Some of the additional factors for consideration are described below.

Single Number vs. Size Range

In November 2024, a subgroup of the NERC LLTF conducted an informal survey to gather feedback from participants on what load size should qualify as "large" under a potential NERC regulatory construct. Most of the survey respondents qualified "large" as greater than 50 MW, and the single size number most commonly suggested was 75 MW.

However, the survey also revealed that "large" may be relative to other factors, such as the voltage to which the load is interconnected. For example, a consuming site with a 20 MW peak demand would be considered "large" for distribution service, whereas 20 MW is a relatively small customer if interconnected at transmission voltage. Additionally, the relative size of the load to the local system and overall interconnection affects how it impacts BPS reliability. For instance, frequency stability looks different in the Eastern Interconnection compared to the ERCOT Interconnection.²⁰ Additionally, even within the ERCOT Interconnection, voltage stability considerations are different in each zone.

¹⁹ Dominion Energy Virginia, "Electric Transmission Facility Interconnection Requirements," Apr. 01, 2025. <https://www.dominionenergy.com/-/media/pdfs/virginia/parallel-generation/facility-connection-requirements.pdf?la=en&rev=07744be57e5a439988e46f630af76e6f> (accessed Jun. 06, 2025).

²⁰ "2024 Frequency Response Annual Analysis," Nov. 2024. [Online]. Available: https://www.nerc.com/comm/OC/Documents/2024_FRAA_Report_Final_Draft.pdf

While defining a singular size could be a way to establish a clear threshold for future regulations, policies, and standards, it could also create unjust or unreasonable unintended consequences. For example, a single number may motivate regulated entities to intentionally size beneath to avoid regulation. This is not to say that size is irrelevant; it is rather that size alone would be too limited a way to examine what is a materially impactful large load.

Behind-the-Meter Loads

The increasing prevalence of behind-the-meter (BTM) large loads is an important issue in the landscape of electric markets, grid reliability, and operations. [Figure 2.5](#) illustrates the difference between front-of-meter (FOM) and BTM configurations. BTM large loads receive their energy directly from co-located generators without the use of the transmission or distribution grid. In certain configurations, the load is served partially through FOM and BTM configurations.




Figure 2.5: FOM vs. BTM Load Illustrations²¹

When a load is served exclusively by a co-located generator, protection mechanisms are typically installed to prevent grid energy from flowing to the load. In these instances, if the co-located generator is out of service, the load will also be disconnected unless it possesses backup generation of its own. [Figure 2.6](#) demonstrates the effects of a generator outage when the load is served exclusively by co-located generation.

²¹ The metering configuration in the BTM configuration example is for illustration purposes only; specific metering design for BTM configurations will depend on the metering practices of the interconnecting utility or utilities.




Figure 2.6: BTM Online vs. On Outage

As BTM large loads continue to develop, their impact on operational visibility, independent system operator (ISO) control, and system stability becomes increasingly influential. Below is a list of characteristics that impact the BPS:

- Operational visibility differs between FOM and BTM loads.** FOM loads are directly integrated into the transmission or distribution system and are usually visible to the interconnecting utility and/or the system operator in the operational environment. In contrast, depending on the utility's requirements, BTM loads may not undergo the same level of impact analysis during the interconnection process. As a result, BTM resources can pose reliability risks if their generation unexpectedly trips off-line and appropriate protection or ride-through systems are not in place. The operational visibility and controllability of BTM loads are limited, making it more challenging to predict system impacts, especially during uncontrolled transitions between on-site generation and grid supply that may arise due to protection system failures.
- Load or generator trips pose risks to system stability.** In cases where co-located generation suddenly trips off-line, the co-located load might not seamlessly transition to grid supply unless a formal interconnection agreement exists with the utility, leading to a sudden interruption of service to this load. If the load unexpectedly shifts to grid supply after the co-located generation fails, the utility supply sees a rapid demand spike, which may stress the transmission system.
- Size and interconnection type matter.** Large BTM loads shifting to the grid during generator outages can disrupt system stability and cause capacity shortfalls or resource adequacy issues. Clear interconnection standards, monitoring, and coordination are needed to prevent unexpected demand swings and ensure reliable operations.

Real-Time Load Behavior

Each category of large load—computational, industrial, and hydrogen—displays real-time operating characteristics that differ from loads previously connected to the BPS. These loads are typically not controllable by the utility or ISO unless they voluntarily enter into a load curtailment or demand-response program. The choice to participate in such programs is usually an individual customer strategy based on an economic evaluation (i.e., customers that have flexible processes may be able to earn revenue from such programs through demand reductions during periods of

high grid stress). Certain control areas allow loads to provide ancillary services, like ERCOT's Controllable Load Resource program.²² These programs are all voluntary.

Many large loads, such as cryptocurrency mining and AI facilities, can cycle their consumption on and off in less than a minute and can ramp from zero to hundreds of MW of power demand over very short time frames, like in the previously mentioned [Figure 2.2](#) and [Figure 2.3](#). As a comparison, data centers that support cloud computing and digital services typically have high load factors and non-conforming behavior, so their consumption is relatively static.

Firm Load vs. Flexible Load

Large loads interconnecting to the grid typically request firm utility service during the interconnection process. This means that the utility is required to provide the transmission infrastructure capable of serving the load's peak demand request at all times.

Concepts have been explored²³ in some areas around mandatory load curtailments by the ISO or utility during stressful grid conditions, which might allow for TPs to assume a lower peak demand for the load and potentially reduce the transmission buildout exclusively needed to support the load. Utility controllability of loads is usually not accounted for at the transmission planning stage. However, many loads are still expected to obtain firm transmission service to ensure reliable electric service delivery, necessitating transmission buildouts to support the load and an increase in energy supply during both normal and emergency system operations.

²² "Load Resource Participation in the ERCOT Markets," *ERCOT*, 2021. <https://www.ercot.com/services/programs/load/laar> (accessed May 30, 2025).

²³ T. Norris, T. Profeta, D. Patino-Echeverri, and A. Cowie-Haskell, "Rethinking Load Growth: Assessing the Potential for Integration of Large Flexible Loads in US Power Systems," Nicholas Institute of Energy, Environment, & Sustainability, 2025. Accessed: May 30, 2025. [Online]. Available: <https://nicholasinstitute.duke.edu/sites/default/files/publications/rethinking-load-growth.pdf>

Chapter 3: Risks to the Bulk Power System

The interconnection of emerging large loads may pose reliability risks to the BPS, and the LLTF has reviewed, categorized, and prioritized these risks. Grid planners and operators need to consider large loads in both the planning and operations domains to ensure that the transmission system is properly planned considering these loads. They must also ensure that the system can be reliably operated in real time with adequate reserves. Large loads can pose numerous risks to the BPS reliability, as discussed below.

Large Load Observability and Data Risks

System operators and planners need data and models about large loads to properly characterize the load's behavior and study potential risks to the BPS. To run steady-state and dynamic simulations, for example, the operator and planner need to know the expected interconnection timelines, peak demand, load behaviors, protection and control settings, and dynamic models for the load. Without this data, operators and planners cannot properly account for the load's effect on the rest of the BPS, potentially leading to poor operating and planning decisions and poor real-time situational awareness.

Currently, developers, owners, and operators of large loads are not required to be NERC registered entities under NERC's Registry Criteria Rules of Procedure. This means that large loads are not required to adhere to Reliability Standards. Operators have already cited multiple events in the Eastern Interconnection (EI) and ERCOT systems that involved large loads in which it was difficult to obtain information from the large load customer, impeding the operator's ability to investigate loss-of-load events. While load currently does not need to register or be involved in the investigation, the magnitude of these loads poses a risk to system stability if the loads behave unexpectedly. Additionally, the control systems of these loads can interact with the grid, exacerbating oscillations and other phenomena. If there is no high-speed recording data available, diagnosing the root cause and fixing it may be difficult or impossible.

In addition to impeding proper planning, the lack of high-resolution measurement data from large loads will hinder system operators' ability to analyze events after they occur. Few loads have phasor measurement units (PMU) or other ways to monitor their behavior to validate facility performance and modeling. Some notable grid events that involve large loads are only able to be analyzed and published in reports because of high-speed recording data at or near a large load.

Many large loads—especially data centers and other computational loads—may shift their computational demand to a different physical facility. These shifts may occur in response to changes in factors such as energy pricing, emission intensity, and currency pricing. However, without knowing the exact causes for these shifts and when they will occur, system operators cannot account for the load response or create accurate forecasts, potentially leading them to use more balancing reserves than expected to handle large load power swings.

Long-Term Planning

Large loads must be accounted for in the transmission and resource adequacy planning horizons to reliably plan and operate the grid. Building a new transmission line can take up to 10 years due to permitting, planning, equipment lead times, and construction. With accurate models and studies, planners can ensure that the system is designed to serve the load and maintain resource adequacy.

Demand Forecasting

Large loads pose specific risks to both the short- and long-term demand forecasts. The risks to BPS reliability related to demand forecasting (1+ years) out include the following:

- Under-forecasting the growth of emerging large loads could pose long-term challenges to resource and transmission adequacy.

- Rapid integration of emerging large loads poses challenges to existing assumptions in resource adequacy and transmission planning processes.

Each load by itself may represent a significant portion of both annual energy consumption and peak demand. In aggregate, they comprise even more. In many jurisdictions, the types and volumes of proposed large loads are unprecedented, making them difficult to incorporate into existing forecasting models. As of April 28, 2025, ERCOT had 136 GW of large load in its interconnection queue with energization dates from 2025 through 2030, with a system that has a historic peak of approximately 85 GW.²⁴

Estimating the amount of load that will make it from interconnection studies to energization is a challenge for grid planners. As stated before, the primary BPS risks to resource adequacy and transmission planning arise from under-forecasting the growth of emerging large loads.

The arrival and timing of large loads is also frequently uncertain. Some companies engage in “location shopping,” exploring and submitting interconnection requests to multiple regions to determine the best location for their operations, creating uncertainty about which projects are firm enough to include in forecasts. Some proposed large loads do not materialize as planned: they may be canceled due to economic and financial challenges or other setbacks. Due to the uncertainty of these large loads materializing, forecasting load for resource adequacy and transmission planning may be much more difficult.

Resource Adequacy

Large loads may also introduce risks to BPS reliability through their effects on resource adequacy. The Virginia Joint Legislative Audit and Review Commission reported that construction of individual data center buildings usually takes 12–18 months.²⁵ In a 2024 report on generator interconnection timelines, Lawrence Berkeley National Laboratory showed that the median duration from interconnection request (IR) to commercial operations date (COD) is nearly 5 years, with the 25th percentile around 40 months and the 75th percentile around 70 months.²⁶ As these loads are magnitudes larger than most loads historically seen on the system and proliferating more quickly than most historical load growth, the need for additional generation may not have been planned for, and current system plans may not be able to meet it.

As a result, demand may outstrip generation supply in the near future. Some of this risk may be mitigated if loads are flexible and demand can be reduced during peak hours.^{23,23} But, as mentioned above, most computational and other large loads are expected to request firm service. If demand leads to a shortfall of generation and all operating reserves—including controllable loads—are exhausted, manual load shed may be needed to stabilize the system.

Performing resource adequacy assessments and resource planning poses a challenge when the amount of load that may materialize is variable or uncertain, especially when compared to forecasting growth of traditional loads. Planners may model additional scenarios to account for the uncertainty of the magnitude and timing of large loads, but this can lengthen timelines and workloads. This can also cause confusion, as one scenario will ultimately need to be used for planning purposes.

If loads connect within the 10-year *Long-Term Reliability Assessment (LTRA)* time frame or five-year adequacy studies, their connections may be unaccounted for when undertaking coordinated planning between regions until the study is updated in the following year.

²⁴ “ERCOT Monthly Operational Overview (April 2025).” Accessed: May 30, 2025. [Online]. Available: <https://www.ercot.com/files/docs/2025/05/15/ERCOT-Monthly-Operational-Overview-April-2025.pdf>

²⁵ “Data Centers in Virginia,” Virginia Joint Legislative Audit and Review Commission, Dec. 2024. Accessed: May 30, 2025. [Online]. Available: <https://jlarc.virginia.gov/pdfs/reports/Rpt598.pdf>

²⁶ “Queued Up: 2024 Edition,” Lawrence Berkeley National Laboratory, Apr. 2024. Accessed: May 30, 2025. [Online]. Available: https://emp.lbl.gov/sites/default/files/2024-04/Queued%20Up%202024%20Edition_R2.pdf

Transmission Adequacy

Transmission adequacy is a well-defined industry concept that is embodied in numerous NERC standards from both the planning and operating perspectives. Large loads can be of a magnitude that must be analyzed well in advance to assess the security of the transmission system and determine whether transmission upgrades are needed to reliably serve the load and meet the expectation of an “adequate level of reliability.”²⁷ Failure to analyze the impact of the new large load can also lead to real-time transmission system restrictions in the amount of load that can be served.

Part of transmission planning is evaluating the planned system for thermal overloads, voltage criteria exceedances (high and low), and system stability performance against various performance criteria. All problematic findings must be remediated prior to energization to prevent real-time operating reliability risks. Without comprehensive pre-energization analysis and configuration planning, large loads pose potential risks to the reliable operation of the BPS. The larger the load’s peak demand, the greater the risk that it will contribute to an inadequate level of reliability.

From a planning and interconnection perspective, large load interconnection designs are analyzed similarly to large generator interconnection designs. The performance of the transmission system with the inclusion of the large load and its connecting facilities can be defined in general terms as follows:

- The transmission system does not experience instability, uncontrolled separation, cascading, or voltage collapse under normal operating conditions and when subject to predefined disturbances.
 - The performance outcomes are the following:
 - Stable frequency and voltage within predefined ranges
 - No instability, uncontrolled separation, cascading, or voltage collapse
- System frequency is maintained within defined parameters under normal operating conditions and when subject to predefined disturbances.
 - The performance outcomes are as follows:
 - Stable frequency within predefined range
 - Facility ratings are respected
 - Frequency oscillations experience adequate damping
- Transmission voltage is maintained within defined parameters under normal operating conditions and when subject to predefined disturbances.
 - The performance outcomes are as follows:
 - Stable voltage within predefined range
 - Facility ratings are respected
 - Voltage oscillations experience adequate damping
- Adverse reliability impacts on the transmission system following low-probability disturbances (e.g., multiple contingences, unplanned and uncontrolled equipment outages, cyber security events, and malicious acts) are maintained at an acceptably low level and, if they occur, are mitigated.
 - The performance outcome is to manage the propagation of frequency deviations, voltage deviations, angular instability, uncontrolled separation, and cascading failures.

²⁷ “RE: Informational Filing on the Definition of ‘Adequate Level of Reliability,’” NERC, May 2013. Available: <https://www.nerc.com/FilingsOrders/us/FERCOrdersRules/ALR%20filing%202013.pdf#search=adequate%20level%20of%20reliability>

- Restoration of the transmission system after major system disturbances that result in blackouts and widespread outages of transmission elements is performed in a coordinated and controlled manner.
 - The performance outcome is to recover the transmission system and restore available resources and load to a stable interconnected operating state.

Bulk Electric System (BES) transmission capability is assessed to determine availability to meet anticipated BES demands during normal operating conditions and when subject to predefined disturbances.²⁸ Additionally, significant demand additions can require a high level of concurrent new transmission buildout as well as many transmission upgrades that require complex and long-term outage coordination evaluations. Delays in completing these new buildouts and projects can result in increased risks to transmission adequacy to serve loads from existing and new generation.

Operations/Balancing

Short-Term Demand Forecasting

Many emerging large loads differ technologically from past industrial loads, leading to significant uncertainty in modeling their behavior and consumption profiles. New loads serve new end uses, such as the parallel training of large language models in AI data centers. This uncertainty results in considerable forecasting errors that could cause system operators to run the risk of under- or over-scheduling, dispatching, or procuring energy and ancillary services. Even after these new loads have been in operation for an extended period, real-time forecasting can remain challenging due to their unpredictable and stochastic nature. For example, high-intensity electrical processing loads like EAFs often experience frequent on-off cycles, causing large, random fluctuations in consumption. Such behavior can lead to dispatch errors, shutting down units prematurely, committing units that are not needed, and generally causing balancing issues. Additionally, large loads are particularly difficult to forecast during extreme conditions, such as during periods of high electricity prices or adverse weather events, amplifying errors in dispatch algorithms when forecasting accuracy is most critical.

Concentrating large loads at single points increases exposure to numerous forecast risks. A single fault could trip the load, from a utility perspective, and amplify the grid disturbance. Moreover, if one of these large loads, such as a data center, is lost and does not return quickly, it can skew the overall forecast and affect transmission station-level forecasts and planning.

Currently, many large loads do not submit real-time or day-ahead operational consumption profiles or plans. This lack of visibility creates forecasting issues, especially for loads with unpredictable consumption patterns, such as data centers or EAFs. Existing short-term forecasting models typically rely on regression-based or autoregressive time series methods that assume past behavior will be replicated in a predictable manner, which fails to account for the dynamic nature of these types of loads.

Load Response to Price and Other Signals

Many large loads may change their demand based on signals like price or emission limits. If the utility is unaware of the triggers that cause the load to change its demand, the utility may not be able to accurately forecast demand. ERCOT analyzed the average percentage of each large load's curtailment when prices go above specific thresholds, and a handful of loads were highly responsive to prices. Conversely, another handful of large loads curtailed less than half of their demand when prices were high. ERCOT notes that a large portion of its large loads at the time of the analysis were cryptocurrency mining, which is known to be more price sensitive when compared to something like cloud data centers that prioritize uptimes over 99.999%. When the system operators are not aware of these triggers, the forecasting quality may be degraded, which could lead to insufficient reserves.

²⁸ Reference NERC's "Adequate Level of Reliability" submittal to FERC: [Link](#)

Balancing and Reserves

In real-time power system operations, there must be enough generation to meet demand while maintaining reserves. Failure to maintain this balance could lead to load shed or a system collapse if there is not enough power or if the reserves cannot handle load movement or unit trips.

Large loads, especially PELs, can shift their consumption in seconds, much quicker than conventional generators can ramp. Quick load ramps—increase or decrease—can stress the system, as generation must be rapidly loaded or unloaded in response. These quick real power demand ramps have been observed in the field during normal operations. For example, [Figure 3.1](#) shows a North American data center ramping down from about 450 MW to about 40 MW within 36 seconds around hour 6. The load’s demand is constant (around 7 MW) for approximately 4 hours. After hour 10, the data center ramps back up to 450 MW over the course of a few minutes.




Figure 3.1: Data Center Load Ramp Down and Up

Forecasting these loads to ensure that the system operator procures the appropriate reserves is also a challenge. Without accurate real-time forecasting methods, the system operator may not procure enough energy reserves in the operating day, leading to energy shortfalls and potential load shedding. Loads in operation have already been shown to be flexible, but the signals that affect demand are not known. Without knowledge of the triggers that cause facilities to shift consumption, it will be impossible to properly account for the load behavior in forecasting.

Large load ramp rates may cause issues with frequency regulation by outstripping the reserves held to regulate frequency. These reserves are provided by on-line spinning resources with headroom to manage a decrease in frequency and units that can further lower their output to manage an increase in frequency. With larger ramp rates, more regulation services will be needed to handle larger swings.

In addition to the magnitude of the ramp rates, the timing matters as well. To handle loads that ramp within seconds or minutes and require fast-acting regulation services, system operators typically procure fast-frequency response services to maintain frequency. The case of ramping down may necessitate units that can rapidly shut down.

The system operator will need information on the large load ramp rates to assess the potential magnitude of the ramping problem. Fast ramping of large loads can cause issues with both voltage and frequency regulation.

If the system operator fails to procure adequate regulation to manage frequency, then load shed or rapid generation curtailment may take place and result in unplanned outages, reducing system security/reliability. In addition to the loss of load, the sudden restoration of a large load could exhaust available balancing reserves, ultimately leading to decreased system frequency and potential frequency instability.

Large rapid changes in demand will also rapidly alter the flow of reactive power in the system, potentially exceeding the ability of existing voltage controls to respond and lead to over- and undervoltage conditions. This could lead to the unplanned outages of generation plants initiated by voltage-ride through protection, load loss related to undervoltage load shedding, or even more widespread outages caused by voltage instability and collapse.

Figure 3.2 shows a load ramp down 298 MW within 25 seconds at a cryptocurrency mining facility in North America. The load was exhibiting real power oscillations with a peak-to-peak amplitude of around 25 MW following a load control issue due to an offsite telecommunication failure. The operator instructed the load to decrease its demand, demonstrating the potential for rapid load ramps.




Figure 3.2: Cryptocurrency Mining Facility Load Oscillation and Ramp Down

Lack of Real-Time Coordination

The lack of coordination between grid operators and large load operators creates multiple risks, including those mentioned above, such as load-ramping coordination.

Risks like these may lead to challenges in controlling the area control error (ACE). For example, AI model training at the xAI Colossus Supercomputer in Memphis, Tennessee—currently the world’s largest AI training cluster²⁹—can change the loading 35–70 MW or more within a minute as the model starts and stops. That change may not be an issue for a single facility, but the aggregate effect across multiple facilities may negatively impact ACE control. These sudden ACE changes may be reflected negatively in the BA Control Performance Standard 1 (CPS1) and BA ACE Limits (BAAL) as required in NERC BAL-001.³⁰

The key characteristics impacting real-time coordination for CPS1 and BAAL are: (1) ramp rate (2) peak demand and (3) load predictability. Large, fast, unpredictable ramps from large loads may cause volatility in ACE and BAAL if a BAS

²⁹ P. Kennedy, “Inside the 100K GPU xAI Colossus Cluster that Supermicro Helped Build for Elon Musk,” *ServeTheHome*, Oct. 28, 2024. <https://www.servethehome.com/inside-100000-nvidia-gpu-xai-colossus-cluster-supermicro-helped-build-for-elon-musk/>

³⁰ [Standard BAL-001-2 – Real Power Balancing Control Performance](#)

generation fleet can't make rapid adjustments. Additionally, if a BAs generation fleet is making rapid adjustments to follow fluctuations in large loads, response reserves could end up being depleted.

There have been documented instances where grid events have unexpectedly impacted data center loads. The NERC *Considering Simultaneous Voltage-Sensitive Load Reductions* incident review¹ details how a 230 kV fault led to customer-initiated simultaneous loss of approximately 1,500 MW of voltage-sensitive load that was not anticipated by the BPS operators. There was a corresponding sudden increase in ACE for the BA when the load tripped. A similar but opposite scenario could occur if a fault at a large load facility adversely impacts nearby generation. The risk to the BPS is unnecessary relay triggering causing potential interactions between the load and generation or the sudden loss of generation and the impacts to ACE control.

Outage Coordination in Operations Planning

Large loads that want to connect within 18 months may also pose risks to outage planning by changing the planning forecast. If the load's connection is approved and there is a resulting increase in load, resource and transmission adequacy margins may be affected. Generation and transmission outages that were approved based on a previous load forecast may need to be shifted or delayed if there is now a resource or transmission adequacy issue based on this increased forecast. This could result in delayed maintenance. Maintenance outages that are taken to avoid sudden tripping (like a failing lightning arrestor) being delayed could increase risks of sudden tripping, which may take longer to repair, leaving the equipment out of service for longer than the original maintenance request. Extended outages on key equipment like generators may affect the system's ability to serve all loads with sufficient operating reserves. Additionally, significant demand additions can require a high level of concurrent generation and pipeline infrastructure development to adequately serve the demand. Delays in completing these new buildouts can result in increased risks to resource adequacy.

BAs must accurately determine their total generation capacity to meet both forecasted and real-time demand. A crucial component of this process is outage coordination. Typically, generators are required to submit their availability, including any deratings, to their BA. This critical information is provided at a minimum in real-time operations and in the near-term horizon. Inconsistent outage coordination between BAs and large load operators can result in inaccurate load forecasts, resulting in either under- or over-committing generation resources. If generation is under-committed, there may not be enough generation to serve all the load. Over-committing generation may cause a different reliability issue related to the mechanical limitations of generation. Thermal units have minimum output levels that can be challenging for them to operate below and may cause undue stress on the unit to come off-line. The risk to the BPS is the potential for inaccurate forecasts and suboptimal generation commitments due to a lack of visibility into large load outages.

Stability

The BPS is planned and operated to be stable as system conditions change and disturbances occur. Potential consequences from power system instability (i.e., a loss of stability) include widespread power outages and permanent damage to BPS equipment. Maintaining stability involves considering a broad range of mechanisms from which instability may arise—[Figure 3.3](#) illustrates a widely used taxonomy of power system stability problems. Large loads, owing to their high power ratings, fast controls, rapidly varying load profiles, and heavy use of active power electronics, are capable of contributing adversely to any of the stability problems shown in [Figure 3.3](#).




Figure 3.3: Taxonomy of Power System Stability

To keep stability problems from adversely affecting BPS reliability, it is generally necessary to identify them proactively and put mitigations in place to ensure conditions conducive to instability are avoided. Power system stability studies are the primary means by which this is accomplished. Thus, many of the stability-related risks associated with large loads arise from the concern that the loads have not been accurately represented in power system stability studies. This is one reason why the widespread tripping of large loads has attracted significant attention—this widespread tripping was not anticipated in power system studies and could indicate unforeseen (and therefore unmitigated) stability problems. In the next white paper, gaps in current practices (e.g., modeling practices) contributing to this concern will be discussed in detail. This white paper, however, is dedicated to identifying ways in which large loads might contribute adversely to stability problems.

Ride-Through

The voltage and frequency ride-through behavior of large loads during disturbances plays a significant role in how they may contribute to instability. Ride-through behavior is primarily defined in terms of how long the load remains connected during a given voltage and/or frequency disturbance. However, changes in the load's real or reactive power consumption during or after the disturbance are important as well (e.g., how much time passes before a disconnected load reconnects, and how quickly does it return to its original consumption). Currently, much attention is directed toward the tendency of many large loads to disconnect during disturbances. Some large loads have internal protection and control systems that will disconnect from the grid during disturbances. For example, some data centers may switch to backup power systems after three transient voltage disturbances within one minute as observed at certain data centers in the EI load transfer event.³¹ The intent behind such systems is usually to ensure the reliability of the large load's process (e.g., serving internet traffic) or protect equipment from damage by switching to a local backup power source.

BPS equipment is designed and operated to meet certain ride-through requirements (e.g., Reliability Standard PRC-024) so that it remains on-line during disturbances and supports the system.³¹ These requirements are necessary because disconnecting BPS equipment during a disturbance can worsen the disturbance's effects and, if enough equipment disconnects, lead to cascading power outages. At a sufficient scale, the disconnection of load introduces similar concerns. The tendency of some loads to disconnect during disturbances is not new. However, the scale of recent load loss events was unexpected and had measurable effects on the BPS. This called into question existing modeling practices and expectations for loads.

³¹ [Standard PRC-024-3](#)

Some large loads have disconnected in large quantities during system disturbances, including the following:

- ERCOT has seen numerous unexpected reductions in consumption or loss of load during disturbances. The majority of losses were with PEL at cryptocurrency mining facilities and oil and gas facilities, which reduced their consumption in response to transmission faults. It was noted that not all facility protections were visible to ERCOT or included in dynamic models. Analysts lacked necessary single-phase high-resolution data as well.
- In the EI, a transmission fault caused the simultaneous loss of approximately 1,500 MW of voltage-sensitive load, primarily from data centers. While circuit breakers at the data center substations did not trip, multiple data centers decreased consumption, switching some of their facility power to backup systems in response to the transient voltage disturbance. **Figure 3.4** shows the loss of data center load in the EI following this fault. The third spike is the third recloser shot at 19:00:39.21 upon which the load tripped.




Figure 3.4: Data Center Load Trip

In subsequent sections, stability issues associated with large and sudden load loss are discussed. Ride-through-related challenges are a focal point in many discussions as they are already having an observable impact on BPS dynamics, but there are other stability risks unrelated to ride-through behavior. Power system instability events are often low-probability, high-impact events that require proactive mitigation. Therefore, it is essential to consider potential stability issues unrelated to ride-through, even if they are not yet observed in operations.

Frequency Stability

BPS frequency is maintained at a near-constant value by closely balancing the energy produced by generation resources with that demanded by the load and losses in power system equipment. This is achieved on timescales varying from seconds to years by different processes. Frequency stability risks are associated with the shortest of these timescales and are closely related to operational risks surrounding balancing and reserves. The stability concern is that very large and sudden (e.g., within seconds) changes in the demand from large loads might exceed the ability of fast-acting generation controls to respond and restore balance between generation and demand. As this capacity is exceeded, the frequency deviation will grow larger. Both loads and generation have limits to the frequency deviation that they will withstand before tripping off-line, which may cause further frequency deviations. Thus, the critical reliability concern is that operators may totally lose their ability to regulate the system frequency and large-scale power outages might follow.

When a large load trips, it causes an instantaneous imbalance in load and generation and frequency deviates from nominal. For example, according to NERC's incident review¹ on July 10, 2024, a lightning arrester failure on a 230 kV transmission line in the EI led to multiple system faults within 82 seconds. These repeated faults resulted in voltage depressions ranging from 0.25 to 0.40 per unit in the affected area. Coinciding with this disturbance, approximately

1,500 MW of load was lost, as demonstrated in [Figure 3.5](#)—not due to utility disconnection but because of customer-side protection and controls. This sudden load loss caused frequency to rise to 60.053 Hz before it stabilized back to 60.0 Hz within four minutes, as seen in [Figure 3.6](#). This incident demonstrates how large-scale load tripping can meaningfully affect BPS frequency. Note that the EI is the largest of the interconnections in terms of inertia and frequency responsive generation capacity—load loss events of a similar size can cause much larger frequency deviations in other interconnections (e.g., a ~235 mHz rise in the ERCOT Interconnection).




Figure 3.5: System Load Chart¹




Figure 3.6: Frequency Plot for Eastern Interconnection 1,500 MW Large Load Loss/Transfer on July 10, 2024 (Provided by University of Tennessee: Knoxville)

Generators are prone to damage from overfrequency conditions as an increase in system frequency corresponds to an increase in the generator's mechanical speed as well. For this reason, BPS generators are equipped with

overfrequency protection. Large load tripping may also result in subsequent generation tripping initiated by overfrequency protection. The tripping of generation will aid frequency stability and mitigate overfrequency conditions, and so the primary concern is not necessarily that frequency would increase indefinitely. Rather, the concern is that the subsequent loss of generation could initiate other reliability issues. Generator overfrequency protection is designed to protect generators from damage during extreme disturbances—it is not designed as a method for fast frequency control or to prevent blackouts. Thus, generation overfrequency tripping that follows large load tripping may cause line overloads, voltage regulation issues, or even a subsequent underfrequency event. All these issues have the potential to cause further operation of protection systems. With each protection system operation, more BPS equipment is taken out of service, and the possibility for cascading outages grows.

The known behaviors of large loads indicate that there is the potential for many gigawatts of large load to be lost nearly simultaneously, leading to the overfrequency issues discussed previously. There are other ways in which large loads could cause frequency stability issues. If many large loads energized at the same time, they could cause significant underfrequency events. This could cause other loads to be tripped off-line via underfrequency load-shedding schemes. If enough large load energizes and is not tripped off-line, the frequency may become low enough that generators trip off-line due to underfrequency protection. If this occurs, frequency instability and widespread outages are a serious possibility as each generator that trips off-line further reduces the system's ability to halt the decline in frequency.

Rotor Angle Stability

Large loads introduce rotor angle stability risks primarily as a result of their potential to cause gigawatt-scale changes in BPS real power flows within a few electrical cycles (50 milliseconds in a 60 Hz system). Of the two forms of rotor angle stability (transient and small signal), these characteristics are more relevant in the context of transient rotor angle stability, which is the primary focus of this section. Rapid changes in real power flows cause synchronous generators to experience power swings, wherein generator power outputs oscillate for a time. The greater the change in real power flows, the greater the magnitude of the resultant power swings. If the power swings become too large, the generators become entirely unable to control their output power or frequency (i.e., they lose synchronism) and must be immediately tripped offline or they will be seriously damaged. The portion of the BPS affected by a particular loss of synchronism event can vary greatly, and consequences range from the outage of a single plant to the islanding of large sections of an interconnection.

There are several ways in which known large load behaviors can adversely affect rotor angle stability:

- **Exacerbated postfault power swings:** BPS faults often cause generators to accelerate rapidly until the fault is cleared. Postfault, the energy consumed by loads provides much of the “braking” force that eventually reverses this acceleration and prevents a loss of synchronism. Thus, if large loads trip off-line during the fault and do not resume consuming energy after the fault is cleared, generators may be unable to halt their acceleration and rotor angle instability may result. This is shown below in [Figure 3.7](#), where a large load is tripped in a positive-sequence dynamic stability simulation. Note that multiple generators lose angular stability in seconds.
- **Unexpected violations of stability-related transfer limits:** If there is too much electrical distance between generators and the loads they serve, the generators may lose synchronism. In terms of system requirements, this manifests as stability-related power transfer limits on transmission lines. This becomes a problem in the context of large loads because the sudden tripping of many large loads can lead to almost instantaneous gigawatt-scale shifts in net power exchange for an area. From a stability perspective, the problem is that the nearby generators, being unable to decrease their real power output quickly enough to account for the loss of load, will begin transferring real power to loads further away. This can lead to sudden and unexpected

violations of stability-related transfer limits, in which case the generator(s) near the large loads will lose synchronism.³²

- Reduced stability margins due to generator reactive power absorption:** Most BPS equipment and most kinds of load consume more reactive power than they produce, and BPS generators supply a significant portion of the BPS's reactive power needs. Some large loads, in contrast with traditional loads, produce reactive power (i.e., they are capacitive loads and operate at a leading power factor). This is because the active power electronics commonly used in LEL operate at roughly unity displacement power factor and are equipped with filters³³ that, while primarily intended to attenuate switching noise and harmonics from the electronics, also produce reactive power. While this can benefit the BPS by improving voltage regulation and reducing transmission system losses, it can worsen rotor angle stability. This is because synchronous generators are more prone to rotor angle instability when absorbing significant amounts of reactive power.³⁴ Unless this excess reactive power production is accounted for in studies, the available stability margin may be overestimated, leading to field events where generators become unstable unexpectedly.




Figure 3.7: Simulation Showing Generators Losing Angular Stability Following Nearby Data Center Load Trip

³² M. Peterson, “Data Center Interconnection Studies & Challenges,” presented at the Large Loads Task Force Meeting and Workshop, Apr. 2025, pp. 53–65. Available: https://www.nerc.com/comm/RSTC/LLTF/LLTF_April_Meeting_&Technical_Workshop_Presentations_.pdf

³³ Some of these filters are separate devices but some are built into the power electronics themselves (i.e., they would not be visible in a single-line diagram).

³⁴ Absorbing reactive power is achieved by reducing the generator's internal voltage. This necessitates operating at a higher power transfer angle to deliver a given quantity of real power, reducing angular stability margin.

The first two issues arise from the ride-through issues discussed previously. The third arises from the unusually high penetration of active power electronics in large loads. In all three cases, shortcomings in stability models are the primary source of risk. All three issues relate to the unique properties of large loads, which may not be captured in existing models or studies. Stability issues cannot be reliably avoided if they cannot be identified, and models and studies are an essential tool for identifying reliability issues.

Rotor angle stability risks are especially relevant for new generation plants being sited near large loads. In some cases, the siting of new generation near large loads is motivated by a need to defer or avoid transmission system upgrades. The same transmission constraints that prevented the large loads from being served by existing generation are likely to adversely affect the generation plant's ability to maintain synchronism with the rest of the BPS.

Apart from risks related to transient rotor angle stability, there does exist the possibility for large loads to affect small-signal rotor stability. Small-signal rotor angle stability problems relate to the generators' ability to successfully damp out small power swings and is the main form of stability that power system stabilizers (PSS) are designed to address. Just as active power electronics devices such as STATCOMs may be used to intentionally act as PSS and improve the damping of power swings, it is possible that the active electronics present in large loads might unintentionally lessen the damping of power swings. The risk of this depends highly on the control design of the active power electronics used in large loads. These can vary significantly by vendor and their inner workings are often protected intellectual property. The risk of small-signal rotor angle stability issues involving a large load are higher if the load is connecting at a location where power swings are more poorly damped.

Voltage Stability

Large loads can affect voltage response and stability, with possible impact on transmission and generation elements, including tripping. While low voltage is generally the focus for voltage collapse, overvoltage issues could result in instability; this has occurred on at least one system.³⁵

Transient (Short-Term) Voltage Stability

Transient voltage stability refers to the ability of the BPS to support system voltages by maintaining adequate dynamic reactive power support following large disturbances. Typically, this time frame captures up to 30 seconds after the disturbance. The ramp rate, peak power consumption, and voltage sensitivity are all major factors in the transient time frame of voltage stability. The larger and faster the difference in real and reactive power consumption is during the transient time frame, the greater the risk to the BPS voltage stability. This could be brought on by a load rapidly ramping up or down. Additionally, this rapid change in power demand could be caused by a planned or unplanned instantaneous tripping of the load. Loads with more sensitivity to low or high voltages will carry increased risk of sudden trips, which can lead to voltage instability.

As demonstrated by an ERCOT voltage stability study,³⁶ risks to BPS voltage stability may occur if enough load trips off-line in response to faults. The primary characteristic of large loads that cause this risk is ride-through behavior. As shown in the NERC *Incident Review on Simultaneous Voltage-Sensitive Load Reductions*,¹ certain fault conditions can cause many electrically close loads to trip or reduce demand.

Mid-Term Voltage Stability

Mid-term voltage stability refers to the ability of the BPS to transition from the transient time frame (less than 30 seconds) to the multiple-minute time frame during which system load and generator response should have stabilized. During this period, numerous factors are in play. Voltage instability arises in this time frame when system reactive demands cannot be met by dynamic and static reactive resources with the applied control schemes (e.g., automatic

³⁵ T. Van Cutsem and R. Maillhot, "Validation of a Fast Voltage Stability Analysis Method on the Hydro-Quebec System," IEEE Trans. on Power Systems, Vol. 12, No. 1, pp. 282-292, February 1997.

³⁶ ERCOT, "Load Loss Threshold Analysis", presented at the ERCOT Large Load Working Group (LLWG) Meeting, May 2025. [Online]. Available: https://www.ercot.com/files/docs/2025/05/19/Large_Load_Loss_Analysis_051625_LLWG.pptx

generator controls, dynamic reactive resource control, controlled shunt devices, and on-load tap changers) and unacceptably low voltage or voltage collapse ensues. Large loads may pose a risk to the mid-term voltage stability of the BPS depending on their ramping behavior, peak load, and overall variability.

Steady State (Long-Term) Voltage Stability

Long-term voltage stability refers to the system's ability to maintain steady voltages once a new operating state is reached, typically well beyond the transient time frame. The magnitude of increase and decrease of real and reactive power informs the amount of risk to voltage stability. The larger the magnitude in change, the greater the risk of steady-state voltage instability. Large loads with the largest peaks and troughs in demand may affect the long-term voltage stability of the system more as compared to smaller loads or less variable loads.

Resonance- and Converter-Driven Stability

The risks associated with large loads and both resonance- and converter-driven stability have many facets in common, and they are discussed together here to avoid unneeded repetition. Both kinds of stability risk involve the closed-loop controls used in large load electronics interacting with other power system equipment. Any closed loop control has limitations regarding the kinds of signals to which it can respond effectively, and the deployment of power electronics in large loads may uncover limitations that were of little or no consequence when the same electronics were used in smaller quantities and for different applications.

For resonance stability, the concern is that large loads will significantly decrease the damping of resonances, or oscillatory modes, in the BPS. These resonances may be electrical (e.g., those introduced by series-compensated transmission lines), mechanical (e.g., turbine torsional modes). Whenever a disturbance occurs in the grid, some oscillation occurs at these resonant frequencies. When the system is stable, the energy associated with the oscillation is absorbed by a mix of generators, loads, and power system equipment, and the oscillation quickly decreases in amplitude. In the context of resonance stability, the concern is that the controls in the large load's electronics will respond by *supplying* energy to the oscillation, instead of absorbing it. Depending on how much energy the large load contributes, the oscillation may increase in duration or even begin to grow in amplitude over time (i.e., the system will become unstable).

The term “converter-driven stability” was introduced to categorize some stability problems that are associated with converter-interfaced generation but not with traditional synchronous generation. Many of the active power electronics devices used in large loads are capable of exhibiting the same converter-driven stability problems as converter-interfaced generation. Converter-driven stability problems can involve interactions between different power electronics and, given the large variety of power electronics designs in use, the possible behaviors and scenarios are many. One form of converter-driven stability that has significantly impacted grid reliability before is weak system stability, an issue associated with the phase-locked loops used by many active power electronics (including those used in large loads). Such controls can become unstable at low short-circuit ratios (e.g., 2 or less). In 2019, an 800 MW offshore wind plant in Great Britain experienced a weak grid instability issue and tripped offline,³⁷ eventually contributing to an outage affecting over 1 million customers.³⁸

Regardless of whether resonance- or converter-driven stability is involved, the potential for the large loads to cause issues is primarily determined by the control algorithms used in the active power electronics. This introduces the same challenge as that associated with small signal rotor angle stability—control algorithms in power electronics vary significantly by make and model, and the control details needed to accurately model the device's potential to interact

³⁷ Y. Cheng *et al.*, “Real-World Subsynchronous Oscillation Events in Power Grids With High Penetrations of Inverter-Based Resources,” in *IEEE Transactions on Power Systems*, vol. 38, no. 1, pp. 316-330, Jan. 2023, doi: 10.1109/TPWRS.2022.3161418.

³⁸ Energy Emergencies Executive Committee, “GB POWER SYSTEM DISRUPTION – 9 AUGUST 2019,” Department for Business, Energy, and Industrial Strategy, Sep. 2019. Available:

https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/836626/20191003_E3C_Interim_Report_into_GB_Power_Disruption.pdf

with grid dynamics are usually protected intellectual property. That said, some risk factors relevant in the case of known stability issues may be relevant for large loads as well.

Some trends from known converter-driven and resonance stability issues might indicate which large loads are at a higher risk of being involved in similar problems. Broadly speaking, the resonance and converter-driven stability issues are more common when a large concentration of power electronics devices is located close to a generation plant (traditional or inverter-based), series compensated line, or other large concentration of power electronics devices (e.g., a STATCOM or HVdc line). The stability issues are usually most prominent when the components involved are more weakly coupled with the rest of the BPS (e.g., they are connected through only a single transmission line). This tends to correspond to a low short-circuit ratio (relative to the size of the large load).

Forced Oscillations

In the context of resonant stability, the concern was that the large load's power electronics cause existing oscillatory modes in the power system to become less damped and potentially unstable. In such cases, the large load is not the source of the oscillation—it is responding to an oscillation that is a natural result of the power system's dynamics. However, there are also cases where large loads can be a source of oscillations and introduce related reliability risks. In such cases, where the large load acts as a relatively fixed source of an oscillatory signal (e.g., voltage magnitude, real power), the load is producing a forced oscillation.³⁹

Broadly speaking, the reliability risks associated with forced oscillations are not always well defined. In some cases, they may cause accelerated aging of BPS equipment (generators in particular) or the operation of protection systems (especially those designed to detect unstable oscillations). In such situations, the concerns generally arise because the frequency of the forced oscillation is the same as that of an existing oscillatory mode in the power system. There are some known risks in such cases, which will be the focus of this section.

If the forced oscillation does not happen to be near the frequency of any known oscillatory modes, it becomes difficult to define specific reliability risks. That said, unexpected interactions involving oscillations have been responsible for well-known reliability issues, such as the Mohave generator shaft failures of the 1970s (caused by subsynchronous resonance, a kind of resonance stability problem) or the 1996 blackouts in the Western Interconnection (which involved an unstable interarea oscillation). The BPS was not designed to operate with large and persistent subsynchronous oscillations, and their presence heightens the risk of unintended interactions that could result in power outages and/or BPS equipment damage.

Some large loads can produce forced oscillations at frequencies below 60 Hz (i.e., subsynchronous frequencies). These oscillations can occur either as a natural byproduct of the load's end use or as the result of unexpected control interactions (i.e., load equipment is not behaving as designed/intended). AI data centers can have load profiles that are periodic, repetitive, and sustained in nature (see [Figure 3.8](#)). More traditional large loads, such as EAFs, can also exhibit sustained subsynchronous oscillations as a natural byproduct of the arcing process.

³⁹ Energy Systems Integration Group's Stability Task Force, "Diagnosis and Mitigation of Observed Oscillations in IBR-Dominant Power Systems," Energy Systems Integration Group, Aug. 2024. Available: <https://www.esig.energy/wp-content/uploads/2024/08/ESIG-Oscillations-Guide-2024.pdf>




Figure 3.8: Example AI Data Center Load Profile During Training Over Two Minutes (Top) and in a Five-Second Period (Bottom)

Forced oscillations can also arise from unintended control issues involving large load equipment. In 2023, a large data center in the Midwest produced forced oscillations unexpectedly. A 1 Hz forced oscillation occurred when a natural frequency was stimulated by periodic forcing of a slower frequency (see [Figure 3.9](#)). The system was inadvertently perturbed at one-second intervals by active power electronics at a data center. Each perturbation resulted in a relatively well-damped 11 Hz ringdown.⁴⁰

⁴⁰ L. Zhu, E. Farantatos, and L. Chen, "Sub-Synchronous Oscillation Detection and Analysis – Dominion Case Study," presented at the Sub-Synchronous Oscillations (SSO) Workshop, Accessed: Jun. 07, 2025. [Online]. Available: <https://www.epri.com/events/539b60d7-57da-4252-9968-fb1754ee3b66>




Figure 3.9: Field Measurement of Forced Oscillations Occurring at 1 Hz⁴⁰

Forced oscillations that occur at modal frequencies present a heightened reliability risk (compared to forced oscillations at other frequencies) because they propagate more effectively—either across the BPS or into specific components such as the blades or shafts of large turbines.

Forced oscillations at interarea modal frequencies can affect an entire interconnection. In 2019, a steam turbine in Florida created a forced oscillation with an amplitude of 200 MW and a frequency of 0.25 Hz.⁴¹ The EI has an interarea mode with a frequency of roughly 0.25 Hz, and this mode is excitable from the Florida region.⁴² Because of this mode, the oscillation spread far across the EI, and oscillations as large as 50 MW were observed in New England. The oscillation persisted for 18 minutes and no outages or damage were reported. However, if the oscillation were larger in magnitude, it may have affected a larger portion of the EI. Depending on the location where the large load is connected and frequency of oscillations it produces, forced oscillations produced by a large load could interact with interarea modes in much the same way and pose risks for a large portion of the interconnection they are connected to.

Forced oscillations at the torsional modal frequencies of a turbine-generating units (steam and gas turbines, specifically) can cause persistent, large torque pulsations (e.g., 30% peak-to-peak) and fatigue the turbine, aging it more rapidly. As an example, [Figure 3.10](#) illustrates the torques produced in a steam turbine by the acceptance of two different prospective EAF loads.⁴³ In the left-hand plot, the EAF characteristic (not shown) does contain significant oscillations, but these occur at frequencies other than the torsional frequencies of the turbine-generating units. Consequently, these oscillations do not propagate very effectively into the turbine and generate little torque ripple beyond an initial transient. In the right-hand plot, however, the EAF's oscillations correspond to one of the turbine-generator's modal frequencies and significant torque pulsations are apparent, even in steady-state.

⁴¹ NERC, "Eastern Interconnection Oscillation Disturbance Forced Oscillation Event," NERC, Jan. 2019. Available: https://www.nerc.com/pa/rrm/ea/Documents/January_11_Oscillation_Event_Report.pdf

⁴² Interarea modes are generally most excitable from locations near the edges of the interconnection and less excitable from regions near the center.

⁴³ [Torsional Resonance Identification in Turbine-Generator Shaft Due to the Operation of Electric Arc Furnaces at Steel Mills in Bangladesh Power System](#)




Figure 3.10: Torques Produced in a Steam Turbine for EAF Loads That Do Not (Left) and Do (Right) Contain Frequency Content at the Steam Turbine’s Torsional Modal Frequency

As with resonance and converter-driven stability, the risks associated with forced oscillations from large loads are higher when the large load is located near a kind of device known to participate in stability problems, as well as when the large load’s short-circuit ratio is low. The same list of devices (generators, capacitors, concentrations of power electronics) is relevant. Forced oscillations involving regional or interarea modes are a special case; studies or measurements must be used to identify the excitability of the modes at buses throughout the BPS. Risks associated with forced oscillations are higher at buses where modes are more excitable,⁴⁴ and modes with lower damping coefficients are more likely to introduce reliability issues if excited.

In the Dominion system, oscillations have been seen arising from the interaction of data centers’ UPS input units.⁴⁵ These oscillations could also interact with nearby inverter-based resources and complicate system stability.

Power Quality

Depending on the operating performance, large loads can have relatively low energy consumption while idling and then have sudden demand spikes when expected to operate. The extensive use of power electronics-based devices could make data centers a significant source of harmonics, unless filtering is designed to address those harmonics. During the transition to these higher-power pulses, the system may experience even higher harmonic distortions, voltage fluctuations causing flicker (visible, frequent changes in the brightness of lights), unbalances, or general power quality issues.

Harmonics

Traditional large industrial loads are known to be a key source of harmonic current injections into utility systems. EAFs employ power-electronic-rich devices such as rectifiers and variable speed drives, which are unbalanced and variable loads. This can lead to significant harmonic current injections, as illustrated in the example harmonic spectrum for an EAF in Table 3.1. In addition, inter-harmonics from these loads have the potential to excite torsional modes of nearby fossil and nuclear generation, which can damage turbine shafts.

⁴⁴ Western Interconnection Modes Review Group, “Modes of Inter-Area Power Oscillations in the Western Interconnection,” WECC, 2021. Accessed: Jun. 09, 2025. [Online]. Available: <https://www.wecc.org/sites/default/files/documents/meeting/2024/Modes%20of%20Inter-Area%20Power%20Oscillations%20in%20the%20WI.pdf>

⁴⁵ Mishra, Chetan & Vanfretti, Luigi & Jr, Jaime & Purcell, T & Jones, Kevin. (2025). Understanding the Inception of 14.7 Hz Oscillations Emerging from a Data Center. 10.13140/RG.2.2.19971.82720. Available: https://www.researchgate.net/publication/389098360_Understanding_the_Inception_of_147_Hz_Oscillations_Emerging_from_a_Data_Center

Table 3.1: Example Harmonic Spectrum of an EAF	
Harmonic Order	Magnitude (% of Fundamental)
2	8.9
3	5.7
4	3.0
5	3.7
7	2.1

Large loads such as data centers are a source of harmonics due to extensive usage of power electronics in both their IT (e.g., UPS, power supplies) and cooling (variable speed drives) components. [Figure 3.11](#) illustrates the excessive voltage harmonics distortion caused by a data center facility and the impact of a harmonic mitigation solution. As shown in the figure, voltage distortion was greatly reduced once harmonic mitigation measures were implemented for the system.




Figure 3.11: Voltage Distortion Before and After Harmonic Correction for a Data Center Facility

Some other large loads may have a variable frequency spectrum that must be accommodated for multiple operating scenarios. [Figure 3.12](#) provides the harmonic spectrum of a medium-frequency induction furnace (MFIF) for both starting and normal operations. Frequency spectrums will often change depending on factors including the operating mode, system configuration, and capacitor switching.




Figure 3.12: MFIF Current Spectrum for Starting and Normal Operation

Voltage Fluctuations

Traditionally, persistent voltage fluctuations resulting in excessive flicker have been a concern with large industrial loads with fluctuating power demand, such as EAFs and welders. [Figure 3.13](#) shows an example demand profile illustrating high variability in power demand of an EAF.




Figure 3.13: Load Profile of an Electric Arc Furnace

Some new large electric loads, such as cryptocurrency miners and AI data centers, also have the potential to introduce significant voltage fluctuations to the supply system. This can be attributed to the significant variation in the load profile of these loads. An example variable load profile of an AI data center is shown in [Figure 3.14](#). This can be an issue, for example, when these loads are concentrated in regions where background flicker is already high (i.e., low

system margin) and these new loads contribute to push the overall flicker above the acceptable values. **Figure 3.15** shows an example of a data center load pulse where the doubling of the load caused voltage distortion and flicker issues for those connected to the facility.




Figure 3.14: AI Training Load Profile Example




Figure 3.15: Data Center Load Pulse (Top: Voltage; Bottom: Line Current)

Voltage Sags and Transients

Voltage sags associated with normal system fault clearing have been known to result in tripping of large load facilities. In addition, transient events associated with capacitor or reactor switching in power systems have also been known to result in unnecessary tripping of large loads, resulting in significant system impact.

Physical and Cyber Security Risks

Cyber Security

This section focuses on theoretical situations and the likelihood of these sorts of attacks are not explored in this paper. The primary cyber security risks to BPS reliability from large loads are bad-actor control of large load demand (e.g. tripping or rapid ramping) and loss of critical communications between a large load and the utility.

It has been shown that some of the latest large loads may exceed 1 GW at a single site. Bad-actor control of a single site that large could affect BPS reliability. If a bad actor has control of the large load, they may be able to control the demand. The risks associated with fast ramping and sudden tripping of large loads are discussed in prior sections of the paper (e.g. Operations/Balancing and Stability).

Loss of communication between the large load and utility could cause similar issues to bad-actor control of the load. If the utility loses visibility of the large load or communications with the large load operators, it could negatively impact operations, balancing, and stability.

Load-Shedding Programs and System Restoration

Manual Load-Shed Obligation Impacts

TOs and TOPs are required to maintain load-shed and under-frequency load shed (UFLS) programs. As large loads are integrated to the system, they may add to the owner's UFLS obligation, eventually reaching a point where the owner must shed all its residential load but still cannot meet load-shed obligations. Tripping the large load could shed more load than intended, risking the stability of the system.

Automatic Under-Frequency Load-Shedding Programs

UFLS programs are the last line of defense for BPS reliability. If frequency declines below preset thresholds, and Transmission Owners have exhausted all preplanned manual load shedding options (i.e., rotating blackout), UFLS programs automatically shed load to arrest the frequency decline and stabilize the system.

Although the PRC-006-5 automatic UFLS standard calls for having a certain amount of load under automatic control to be shed, the rapid addition of large loads makes it even more important to ensure that the automatic UFLS programs are up to date and can address the presence of the new large loads on the system. The current procedure of evaluating UFLS programs every five years may not be often enough. These large loads present a risk that the local distribution system may not contain enough load to cover the UFLS requirements. TOs may need to start including load at the transmission system to meet the requirements of PRC-006.

The functionality of this "last-resort system preservation" program is assessed through studies that identify the electrical islands that may be formed under simulated conditions. The studies are used to establish the parameters of the UFLS entity automatic UFLS programs as required by the standard.

System Restoration

System restoration (blackstart) is the worst-case scenario for everyone who relies on the electric grid, especially for system operators who must restore the grid as quickly and safely as possible. Fortunately, system operators have existing procedures and train annually for this scenario.

Blackstart is an iterative process that relies on either starting a blackstart resource or building out from intact islands to re-energize transmission lines and restore load while maintaining frequency and voltage within acceptable limits. When starting from blackstart resources, transmission lines and loads are brought online iteratively, creating a small electrical island. Multiple islands are gradually connected until the entire grid is re-established. Appropriate load pickup is critical to successful restoration from a blackstart resource because small islands are more sensitive to small changes creating larger swings in voltage and frequency. It is crucial to consider the effects of cold load pickup during this process. If frequency swings too low during the load pickup, it could cause the blackstart resource or other generators to trip. The blackstart resource and other generators should be dispatched near the middle of their real power capability range to allow room to increase or decrease generation to control frequency. Additionally, high voltages could also lead to generators tripping, causing the island to collapse. A summarized example blackstart procedure is listed below.

1. First, small amounts of load (10–15 MW) are restored in a particular island.
2. As islands are stabilized (more transmission lines, more load, more units connected in the same island), the operator reviews the feasibility of connecting two islands together.
3. If connection is successful, more transmission lines, load, and units are added to the now-growing island. If not, the smaller islands must be rebuilt.
4. An island may black out even prior to connecting with another island. The most likely cause in those cases is the addition of too much load at one time—a distinct possibility in the world of new large loads.

System operators prioritize the restoration of loads based on their function, size, and location. Loads are restored in blocks that are limited by the size of the blackstart generator on isochronous control and, eventually, constant frequency control.

Two other load characteristics can affect the system restoration process as loads are restored: segmentation and demand variability. In the case of large loads, the total load of a customer may need to be portioned into smaller, manageable and predictable segments for restoration. A lack of clear, pre-determined segmentation agreed between load customer and TSP during the restoration state can cause frequency decline or voltage collapse. This can trigger UFLS and crash the island or create the need for load shed. The same is true when a restored load increases its demand without system operators' instruction.

Typically, the system operator has had finer control of load segmentation on the distribution and transmission system. However, large loads with internal segmentation raise the risk of restoring too much load too fast.

Chapter 4: Conclusion

As emerging large loads seek to swiftly interconnect to the BPS, they pose new risks to the BPS. Many current large load constructs are based on transmission facility ratings with no additional accounting for the potential impact that large loads may have on the BPS. If large loads are defined by a single number, many in the sector would choose 50 or 75 MW. However, additional characteristics should be considered in any definition of large loads. This white paper shows that in addition to peak demand, many other characteristics of emerging large loads affect their impact on BPS reliability. These characteristics include fast interconnection timelines, demand profile, load predictability, ramp rate, PELs, voltage sensitivity, and internal segmentation. In addition to the characteristics of the load, the system characteristics also affect what a “large load” means. Frequency stability depends heavily upon interconnection-wide properties like inertia and on-line generation. These interconnection-wide characteristics vary greatly in North America. Additionally, risks like voltage instability may be more dependent on local system properties, like available reactive power from nearby equipment.

Large loads pose risks in both the planning and operations horizons. Their quick interconnection timelines and large peak demands drive generation and transmission adequacy risks. The fast ramp rates and variability of the loads could exhaust reserves for balancing and contribute to voltage and frequency instability. If system planners and operators lack accurate dynamic models, they may be unable to predict ride-through and system behaviors during events. Much of the PEL load can contribute to harmonics and voltage fluctuations. Large loads like data centers can also be susceptible to cyber-attacks that could trigger load ramps. Additionally, the rapid pace of load integration with large loads’ magnitude could negatively impact system resilience. Large loads must be considered in load-shed obligations, UFLS program design, and blackstart restoration.

To understand how these risks align with the unique characteristics of large loads discussed in this white paper, see [Appendix B: Comparison of Characteristics and Risks](#).

Given the multitude of risks posed by large loads, it is recommended that the risks be prioritized in the order shown across the following three tables. NERC’s Framework to Address Known and Emerging Reliability and Security Risks states that prioritization of risks is accomplished through analysis of their exposure, scope, and duration as well as impact and likelihood.

High-Priority Risks	
Long-Term Planning	Resource Adequacy
Operations/Balancing	Balancing and Reserves
Stability	Ride-through
	Voltage Stability
	Angular Stability
	Oscillations

Medium-Priority Risks	
Operations/Balancing	Short-Term Demand Forecasting
	Lack of Real-Time Coordination
Long-Term Planning	Demand Forecasting
	Transmission Adequacy
Stability	Frequency Stability
Security Risks	Cyber Security
Load Shedding Programs & System Restoration	Manual Load-Shed Obligations
	Automatic UFLS Programs

Low-Priority Risks	
Power Quality	Harmonics
	Voltage Fluctuations
Load Shedding Programs & System Restoration	System Restoration

Finally, NERC would like to formally thank all the industry experts, from utilities, regulators, national labs, government agencies, and data center developers/operators/owners, who contributed to this white paper. The names of individual contributors and their organizations are provided in [Appendix C: Acknowledgements](#).

Appendix A: Large Load Construct Data

Table A.1: Summary of Large Load Constructs

Region/Entity	Threshold	Rationale	Classification
ERCOT	75 MW	Above 75 MW, a transmission upgrade is likely needed to serve the full load.	Peak Demand
NYISO	10 MW at 115kV or above, or 80 MW below 115kV	These loads could impact the New York state transmission system and need to be evaluated to determine their responsibility for upgrades required to reliably interconnect to the NY State Transmission System.	Peak Demand and Voltage Level
Dominion	100 MW	Above 100 MW, a ring bus configuration is needed per Facility Interconnection Requirements.	Peak Demand
Grant PUD	2 MVA for large, and 40 MVA	2 MVA is the largest secondary service transformer and 40 MVA is the largest transformer size available and requires additional or new substation work. Above 40 MVA may require transmission service.	Peak Demand
Portland General Electric	1 MW, 30 MW	The primary drivers are rate and cost allocation. At 1 MW or greater, there can potentially be an impact on feeder mainline or substation vs. a simple line extension for smaller load. Load at 30 MW aligns with the typical distribution substation transformer loading criteria.	Peak Demand
PacifiCorp	1 MW	Currently driven by tariffs.	Peak Demand
SRP	10 MW, 150 MW	At 10 MW or greater a dedicated substation off of 69 kV system is required. At 150 MW or above, a 230 kV connection is needed.	Peak Demand

Appendix B: Comparison of Characteristics and Risks

Table B.1 shows a mapping of potential large load characteristics to the risk that it causes or is related to. Large loads can be categorized, as shown in [Chapter 2](#). It should be noted that each large load may have a unique mix of characteristics. Additionally, each characteristic may have different thresholds for causing risks depending on the system and local area it connects to.

Table B.1: Large Load Characteristics and Risks Mapping									
Risks	Characteristics								
	Peak Demand	Fast Interconnection Timelines	Demand Profile	Load Predictability	Ramp Rate	Load Type (PEL)	Voltage Sensitivity (Ride-through)	Inaccurate Dynamic Models	Internal Segmentation
Inaccurate Long-Term Forecasts	X	X		X					
Steady State Thermal and Voltage Violation	X		X	X	X		X		
Insufficient Generation Adequacy	X	X		X					
Inaccurate Short-Term Forecasts			X	X	X	X	X		
Balancing Reserves Shortage		X	X	X	X	X	X		
Voltage Instability	X		X		X	X	X	X	
Frequency Instability	X				X		X	X	
Oscillations			X	X		X		X	
Harmonics			X			X			
Insufficient UFLS Procurement	X	X							X
Insufficient Load Shed Obligations	X	X		X					
Loss of Blackstart Islands				X			X		X

Appendix C: Acknowledgements

This white paper was developed by a significant number of contributors from the NERC LLTF. The LLTF would like to acknowledge and thank the following individuals for their support in drafting this white paper.

Table C.1: Contributors

Name	Company
Meg Albright	Bonneville Power Administration (BPA)
Rahul Anilkumar	Quanta Technology
Bob Arritt	Electric Power Research Institute (EPRI)
Sagnik Basumallik	New York Power Authority (NYPA)
Julie Booth	North American Electric Reliability Corporation (NERC)
Alexander Carlson	North American Electric Reliability Corporation (NERC)
Valerie Carter-Ridley	North American Electric Reliability Corporation (NERC)
Candice Castaneda	North American Electric Reliability Corporation (NERC)
Horea Catanase	Electric Power Engineers (EPE)
Michael Cohen	MITRE
Amin Dadashzade	Zero-Emission Grid LLC
Sudipta Dutta	Electric Power Research Institute (EPRI)
Zia Emin	Electric Power Research Institute (EPRI)
Arman Ghasaei	EdgeTunePower
Jack Gibfried	North American Electric Reliability Corporation (NERC)
Patrick Gravois	Electric Reliability Council of Texas (ERCOT)
Martha Henson (lead)	Oncor
Sam Holeman	Duke Energy
James Huang	Enchanted Rock
Pranav Kishore	1898 & Co.
Adam Kliber	Independent Electricity System Operator (IESO)
Brett Kruse	Calpine
Mike Mathieu	Elevate Energy Consulting
Eric Meier	Electric Reliability Council of Texas (ERCOT)
Evan Mickelson	North American Electric Reliability Corporation (NERC)
Shivani Nathoo	Independent Electricity System Operator (IESO)
Brett Ross	Pacific Northwest National Laboratory (PNNL)
Herb Schrayshuen	Power Advisors, LLC
Mostafa Sedighizadeh	Southwest Power Pool (SPP)
Mohamed Shamseldein	Independent Electricity System Operator (IESO)
Harish Sharma	Southern Company
Agee Springer (LLTF vice-chair)	Electric Reliability Council of Texas (ERCOT)
Tyler Springer	American Electric Power (AEP)
Nathan Strobel	North American Electric Reliability Corporation (NERC)
Kyle Thomas (lead)	Elevate Energy Consulting
John Tomczak	-
Matthew Veith (LLTF chair)	American Electric Power (AEP)
Scott Warren	Southern Company
Matthew Wharton	PJM
Jiecheng Zhao	Elevate Energy Consulting