

Characterizing Large Loads

A Taxonomy to Support Large Load
Integration

OCTOBER 2025

Peter Fitzmaurice
Jonathan Tacke
Ryan Cobey
Layne Nelson

Idaho National Laboratory

INL/RPT-25-88356
Revision 0

Grid Deployment Office



DISCLAIMER

This information was prepared as an account of work sponsored by an agency of the U.S. Government. Neither the U.S. Government nor any agency thereof, nor any of their employees, makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness, of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. References herein to any specific commercial product, process, or service by trade name, trade mark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the U.S. Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the U.S. Government or any agency thereof.

Error! Reference source not found.

Revision 0 [only after initial release]

Characterizing Large Loads

A Taxonomy to Support Large Load Integration

**Peter Fitzmaurice
Jonathan Tacke
Ryan Cobey
Layne Nelson
Idaho National Laboratory**

october 2025

**Idaho National Laboratory
Center for Securing Digital Energy Technology
Idaho Falls, Idaho 83415**

<http://www.inl.gov>

**Prepared for the
U.S. Department of Energy
Grid Deployment Office
Under DOE Idaho Operations Office
Contract DE-AC07-05ID14517**

Page intentionally left blank

ABSTRACT

The United States power industry is currently facing unprecedented challenges due to the reshoring of American industry and the rapid advancement of artificial intelligence (AI) technology. As these trends persist, it is becoming increasingly crucial to identify methods for the rapid integration of high-demand electricity customers, such as industrial facilities and data centers. This paper begins to address this challenge by presenting a common approach to characterize these “large loads” through the development of a shared taxonomy.

This work begins with an examination of what makes a particular load “large” and why these loads are becoming increasingly relevant. Next, the essential characteristics of large loads are identified and used to categorize them into types. The paper concludes with a demonstration of the classification taxonomy using several hypothetical examples. This work provides a foundation for utilities, grid operators, and other stakeholders to address the challenges and opportunities related to grid resilience and stability posed by integrating diverse large loads.

It is acknowledged that there is nuance that differentiates large loads (including data centers), particularly concerning operational characteristics and reliability needs. This work is about the challenges to the stability of the grid by the integration of large loads. Therefore, the characteristics that define how a large load interacts with the stability of the power system are purely operational. It should also be noted that this taxonomy and the discussion surrounding it within this paper are intentionally basic in nature and do not delve into various technical and more nuanced topics that are undeniably important to the issue of large load behaviors and grid stability. These relevant issues will be integrated into subsequent work, as discussed at the end of this paper.

Page intentionally left blank

CONTENTS

ABSTRACT.....	v
ACRONYMS.....	ix
1. INTRODUCTION.....	1
2. WHEN IS A LOAD “LARGE”?.....	2
3. KEY CHARACTERISTICS FOR DELINEATING LARGE LOADS.....	3
4. LARGE LOAD TAXONOMY.....	7
5. APPLYING THE TAXONOMY.....	8
6. CONCLUSION.....	13
7. REFERENCES.....	15

FIGURES

Figure 1. Large load landscape.....	3
Figure 2. Differing load profiles.....	6
Figure 3. Large load taxonomy.....	8
Figure 4. Large load type A-2 example.....	10
Figure 5. Large load type B-1 example.....	10
Figure 6. Large load type D-2 example.....	11
Figure 7. Large load type B-2 example.....	12
Figure 8. Large load type C-3 example.....	12
Figure 9. Large load type A-2 example.....	13

TABLES

Table 1. Varying large load thresholds.....	2
Table 2. Three key characteristics.....	4
Table 3. Load size.....	5
Table 4. Slow vs. fast loads.....	6
Table 5. Flexible vs. inflexible loads.....	7

Page intentionally left blank

ACRONYMS

AI	Artificial Intelligence
ERCOT	Electric Reliability Council of Texas
GDO	Grid Deployment Office
ISO	Independent System Operator
IT	Information Technology
MISO	Midwest Independent System Operator
MW	megawatt
NERC	North American Electric Reliability Corporation
NYISO	New York Independent System Operator
PJM	Pennsylvania-New Jersey-Maryland Interconnection
POI	Point of Interconnection
RTO	Regional Transmission Organization
UPS	uninterruptible power supply

Page intentionally left blank

Characterizing Large Loads

A Taxonomy to Support Large Load Integration

1. INTRODUCTION

The United States (U.S.) electrical grid has always faced the introduction of ever larger individual loads, as well as the corresponding need for greater electrical supply and stability measures. As the scale of industrial operations and technological development has expanded, these high-demand electricity customers have come to be termed “large loads.” Grid operators and utilities consistently adapted and evolved planning and operational practices to serve these customers of increasing size, introducing advanced load forecasting, demand-side management programs, real-time monitoring, and more. Today, however, large loads are proliferating more rapidly, across wider geographic areas, and reaching extraordinarily high levels of energy demand. In particular, data centers have emerged as technological and economic pillars that must increasingly be built-out across the nation and require high levels of on-demand electricity. Generation capacity must expand to meet the growing demand, a challenge exacerbated by plans in many regions to retire aging generation plants. Furthermore, the operational nature of many of these large loads, especially data centers, is very different than traditional industrial loads, presenting new challenges but also potential opportunities for the stability of the grid [1].

Despite the significant impact of large loads on the grid, there are no nationally standardized definitions or categorizations for these unique electricity customers; rather, classification varies regionally and by organization. To facilitate greater levels of integration, there are now increasing calls for a common approach to characterize large loads and understand their behaviors and attributes. Notably, the North American Electric Reliability Corporation (NERC) has created a Large Load Task Force, which recently published a white paper that discusses the characteristics and risks of large loads [2]. This document begins to detail differences amongst these types of loads, but emphasizes the need for greater examination moving forward, including “a framework for classifying large loads.”

The goal of this paper, while not a direct response to NERC, is to address this noted gap in characterizing large loads, starting with a discussion of why greater definition and delineation of large loads is needed to support grid stability and resiliency. This work then touches on the question of what makes a load “large.” The main contribution, however, is the identification of the essential characteristics that distinguish categories of large loads across a diverse range of facility types and operational models. It is acknowledged that there is nuance that differentiates large loads (including data centers), particularly concerning operational characteristics and reliability needs. This work is about the challenges to the stability of the grid by the integration of large loads. Therefore, the characteristics that define how a large load interacts with the stability of the power system are purely operational. This paper concludes with a basic taxonomy that proposes a set of classifications of large loads with similar characteristics. This taxonomy is intended for all stakeholders involved in data center integration, but primarily considers the perspective of a grid operator. The intent is to focus on challenges of integrating large loads into existing power systems and providing common language so that the setting of expectations on how large loads will be allowed and expected to participate in the bulk energy system can be coordinated across stakeholders.

The taxonomy approach presented here uses a net impact perspective and considers how a grid operator would perceive loads based on their dynamics and other characteristics. Therefore, important variables for individual load operation, such as on-site generation, storage and operational criticality, are factored into the taxonomy implicitly rather than explicitly. We note that there are resilience and

operational considerations that cannot be covered by this taxonomy, but which would be addressed by building on this taxonomy.

The work presented here can serve as a foundation for utilities, grid operators, regulators, and large electricity consumers to confront both challenges and opportunities related to the resilience and stability of the grid posed by the integration of diverse large loads. It establishes a framework within which discussions about interconnection requirements and ancillary services can take place. Therefore, the intended audience cuts across all large load stakeholders. This taxonomy gives utilities a reference from which to approach requirements for different large load types. Grid operators and regulators can utilize the taxonomy to develop strategies and approaches for large load integration and market structures. Data centers and other large loads can turn to the taxonomy for a simple mechanism for understanding decisions and actions they can take to facilitate interconnection and approvals, as well as potentially expand revenues through participation in grid services.

2. WHEN IS A LOAD “LARGE”?

Historically there have not been rigidly applied standards for identifying large loads. Utilities, grid operators, balancing authorities, public utility commissions, federal regulators and loads themselves often approach the question independently and adjust their definitions based on their unique circumstances. However, at the most fundamental level, a large load is an electricity customer whose peak demand can impose new operational costs to the grid and/or trigger new system-wide stability and reliability risks that require upfront assessment and resolution, often involving joint action by the grid operator and the load. Sufficiently large loads may necessitate procurement of stability reserves to account for their size and behaviors, triggering additional operational costs.

While official thresholds vary across entities, the electric power industry tends to view a load as large once it enters the 50-100 megawatt (MW) range. Table 1 provides some examples.

Table 1. Varying large load thresholds.

Organization	Type	Definition
Electric Reliability Council of Texas (ERCOT)	Independent System Operator (ISO)/ Regional Transmission Organization (RTO)	75 MW [2]
New York Independent System Operator (NYISO)	ISO	80 MW [2]
Dominion Energy	Utility	100 MW [2]
North American Electric Reliability Corporation (NERC)	Federal Regulatory Authority	Varies ^a
NERC Large Load Task Force Informal Survey	Various	50 MW [2]
Midcontinent Independent System Operator (MISO)	ISO	Varies [4]
CPS Energy	Utility	40 MW [5]

^a NERC views a large load as “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.” [2]

Given that the main significance of the large load designation in the context of utility operations relates to how these loads can operationally affect the grid, any definitional threshold cannot be a single measure applied uniformly in all cases and for all grids. As noted by NERC, large load designations must account for a range of operational and technological variations, including the nature and size of the grid at issue [2]. For example, while the Electric Reliability Council of Texas (ERCOT) considers 75 MW a large load for stand-alone facilities, it also designates facilities at the 20 MW level as large loads if they are co-located with generation [5]. Voltage levels, points of interconnection (POIs), load clusters, and other variables can affect what is considered large. Defining a large load could even vary within a given utility’s own service area, with different zones having different dynamics and limitations that call for different large load thresholds.

Despite these and other important qualifiers, for the purposes of this report the term “large load” will mean any load that consists of a single facility capable of consuming at least 50 MW of real power. This is because, for most interconnections, a sudden change in load of 50 MW or more cannot be accommodated through hourly energy markets alone, meaning these loads may represent stability risks. Figure 1 provides a generalized view of where some different operational loads and definitional thresholds lie along a MW continuum.




Figure 1. Large load landscape.

3. KEY CHARACTERISTICS FOR DELINEATING LARGE LOADS

Having established a threshold for what qualifies as a large load, there are a wide range of facilities that fit this criterion, from industrial metal smelters to hydrogen production plants to artificial intelligence (AI) data centers. While two large loads may have the same MW peak demand, they can be significantly different in their operational characteristics and load profile. Therefore, they can also vary in the manner and degree to which they introduce instability risks but also grid service opportunities. It is therefore important to understand distinctions amongst large loads that can have critical importance to the operation and stability of the grid.

There are many factors and attributes that may be considered when delineating large loads. Some examples include a load’s specific operational purpose, equipment type, up-time requirement, components of total load, and resistive vs. inductive attributes. Examining these and other factors can advance understanding of how a load performs and interacts with the grid.

However, developing a taxonomy that is generally applicable to all large loads regardless of facility type requires a more fundamental view of distinguishing factors. Such a view takes the more basic perspective of how a load appears to the external grid entity serving it. Therefore, this simplifies the large load classification process, requiring only those characteristics that are most essential for determining a load’s grid stability effects. These include (1) load size, (2) load variability and (3) load flexibility, each of which is defined in Table 2 and discussed in greater detail below.^b After this initial step of basic large load classification, some of the other distinguishing factors can come into greater focus, including as part of examining the operational potential of different loads for providing grid services.

Table 2. Three key characteristics.

Load Size	Load Size, or peak load, refers to the maximum instantaneous power that a load may draw at any point in time.
Load Variability	Load Variability refers to the way a given load consumes energy. Specifically, it relates to how rapidly and significantly a load can vary the amount of power it consumes from moment to moment.
Load Flexibility	Load Flexibility is a measure of how much a load’s power consumption can be altered to accommodate the stability needs of the grid without adversely affecting the essential function of the load.

CHARACTERISTIC #1: LOAD SIZE

Load size serves as a simple and fundamental quantitative differentiator amongst large loads. As depicted in Figure 1, large loads can extend from tens of MWs to hundreds of MWs, with gigawatt-sized loads on the horizon [6]. Load size affects potential stability risks but can also distinguish a load with outsized ability to provide valuable ancillary services. The larger the load, the greater its operation affects the system, and the greater potential instability effects are associated with it [7]. At the same time, however, these significant loads can offer larger-scale stability services to the grid. Nameplate capacity is used to measure the size of a load, but there are nuances to using this single metric for size. Although most loads will rarely operate at their full nameplate capacity, and many large loads take several years to build up to this full capacity, this metric provides consideration for the maximum impact the load could have to the connected system, which allows for conservative stability analysis. NERC notes that other factors, such as interconnection voltage and size of the load relative to the local system and overall connection could also affect the impact the load has on the system [2]. For the purposes of this taxonomy, which attempts to start addressing these challenges in definition with basic building blocks, Table 3 divides large load sizes into three groups.

^b The NERC white paper [2] also highlights these additional characteristics of variability (i.e., “ramp rates” and “real-time load behavior”) and flexibility.

Table 3. Load size.

Size 1: 50-99 MW	These loads typically include large commercial facilities, data centers, and mid-sized industrial operations. While substantial, they generally integrate into existing grid infrastructure with moderate upgrades. Their size allows for relatively flexible siting and interconnection, making them a common entry point for load growth in both urban and rural areas.
Size 2: 100-999 MW	This category encompasses major industrial facilities such as aluminum smelters, steel mills, and large-scale hydrogen production plants. These loads require substantial grid infrastructure and may require dedicated transmission upgrades or new substations.
Size 3: 1,000+ MW	These are rare but transformative loads, including multi-gigawatt data center campuses, direct air capture facilities, or large-scale electrification of industrial clusters. Their integration can reshape regional transmission planning, resource adequacy, and even market dynamics.

These thresholds are grounded in practical distinctions used across utility planning, regulatory frameworks, and grid operations. The 50 MW mark, as discussed previously, often delineates the boundary between standard and large interconnection procedures, making it a logical lower bound for categorizing impactful loads. At 100 MW, loads begin to trigger more rigorous transmission planning and reliability assessments due to their potential to influence regional power flows and system stability. The 1,000 MW threshold represents a qualitative shift. Loads of this magnitude rival utility-scale generation in their grid impact, often requiring custom infrastructure, long-term planning, and policy-level coordination.

CHARACTERISTIC #2: LOAD VARIABILITY

Load variability reflects an electricity customer’s power use behavior, especially as it relates to the degree and consistency of power demand. Some loads have a relatively “flat” and predictable power demand, even if that demand is high. This type of load variability is representative of steady, sustained operations with limited ups and downs in power demand. Variations in power use will occur, but they are well-known or predictable and do not introduce stability considerations.

Alternatively, a load’s power usage can fluctuate significantly and rapidly, and do so in ways that are unforeseen by utilities and balancing authorities. This can include sudden large spikes and significant drops in power usage. Figure 2 depicts stark differences in load behavior, demonstrating that facilities with similar average utilization rates can demonstrate widely divergent variability.

Load profiles for large electricity customers can be important to stability because significant and rapid changes can cause unacceptable voltage or frequency deviations. However, some fast-acting loads have the potential to support the grid if they can be controlled. Assets with the capability to rapidly change their output in response to grid conditions (e.g. flexible, but variable) might be leveraged to provide ancillary services. This concept is discussed further under the next characteristic.

For these assets to provide ancillary services, markets must be established in a way that facilitates participation by large loads. Alternatively, in regions without ancillary service markets, requirements must be established for these loads to provide grid support services as part of interconnection agreements. Either way, infrastructure must be installed to enable communication of dispatch signals from the balancing authority to the large load asset. These services must be procured over different timeframes, with faster services being more challenging to provide and thus more valuable.

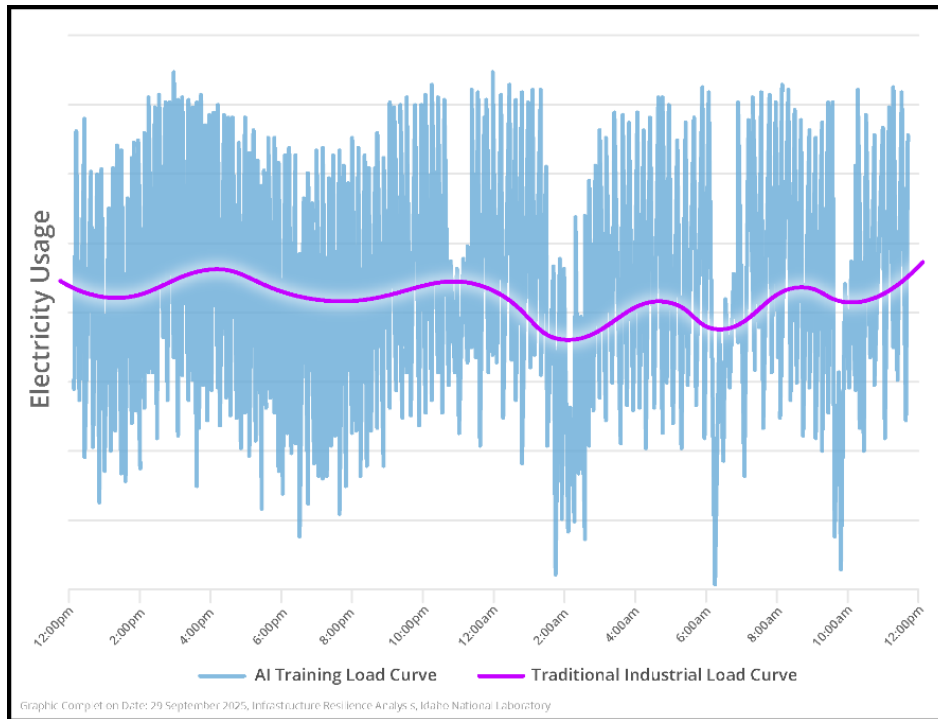


Figure 2. Differing load profiles.

The taxonomy captures these important operational differences by identifying a load as having either a slow or fast load profile, as described in Table 4. The characteristic of load variability starts to address NERC’s call for definition of loads based on real-time load behavior [2].

Table 4. Slow vs. fast loads.

<p>Slow Load Profile</p>	<p>A load with consistent and predictable power demand that rarely experiences significant increases and/or decreases in power usage over short periods. Such a load can be said to vary slowly. An example is a cloud storage data center that seeks to maintain a very level load profile for 24×7 operations.</p> <p>For purposes of this Taxonomy, a “significant” power increase or decrease is equal to at least 50 MW, and a “short period” is considered to be 60 minutes or less. “Rare” occurrences of these rapid power swings are considered to be less than 12 times per year or only during grid instability events. In practice, these thresholds will vary by interconnection^c, but in general these values are those that would spur a balancing authority to take remedial action.</p>
<p>Fast Load Profile</p>	<p>A load that, by virtue of its business operations, experiences significant increases and/or decreases in power usage over short periods. An example is a steel mill with arc furnaces that can widely vary their heat production function during the course of a production run.</p>

^c According to FERC Order 842 (RM16-6), the maximum allowable frequency deadband for primary frequency response is ± 0.036 Hz (0.06% of 60 Hz). If system frequency deviates beyond this range, generators must respond. Since frequency deviation is proportional to the imbalance between generation and load, a simultaneous 0.06% change in all loads across an interconnection would trigger a system-wide generator response within one second. This implies that a single load varying by 0.06% of the interconnection’s total capacity can have the same effect. Therefore, any load that fluctuates by more than this threshold—especially faster than the spot market’s typical 1-hour clearing time—requires additional reserves and is considered “variable.” To simplify, we assume a 50 MW threshold already accounts for interconnection size. Thus, any load capable of changing by more than 50 MW per hour under normal conditions is classified as “fast-moving” in our taxonomy.

CHARACTERISTIC #3: FLEXIBILITY

Similar to Load Variability, Load Flexibility also relates to increases and decreases in power usage, but from the standpoint of voluntary action to support grid stability. This can involve both ramping down and ramping up over defined periods in response to grid needs. Evaluation of a given load’s flexibility characteristics includes any behind-the-meter resources that contribute to the load’s profile at the POI as seen by grid operators. The definition of flexibility in this taxonomy combines considerations from NERC to classify loads based on firm or flexible behavior and consideration of behind-the-meter co-located generators [2].

A load can be considered flexible if it can alter its power consumption from the grid in a controlled manner. The key is that a flexible load can perform these power demand changes (either up or down) without harming its essential operational and business functions, and in response to requests from the grid operator. This characteristic is of interest in grid stability because it determines whether a given load represents a potential asset or liability for balancing authorities. Some key factors that support flexibility include the presence of adequate on-site power generation^d, single tenancy^e, and the ability to work around issues of criticality. Criticality generally involves operations that are essential for human health, security, and national and military missions. While the taxonomy does not have an explicit criticality factor, this important issue is partially incorporated through the Flexibility characteristic. For example, a load may be less likely to enroll in voluntary demand response programs because of its criticality, even if it has on-site power it could utilize, meaning that it would be less flexible from the grid operator’s perspective. Flexibility includes both the technical capability to rapidly adjust energy consumption as well as the business alignment to do so if called upon.

Table 5. Flexible vs. inflexible loads.

Flexible Load	A load that can automatically raise or lower its net power consumption based on requests from the grid (i.e., ancillary services), while avoiding its own operational degradation.
Inflexible Load	A load that is unable to adjust power consumption quickly for purposes of ancillary services due to resultant operational impacts. This inability can be driven by factors such as lack of on-site storage or generation, lack of grid communications capability, mission criticality, and/or organizational rules and policies.

4. LARGE LOAD TAXONOMY

While many factors can influence how large loads are distinguished, the three fundamental characteristics presented above allow for a basic taxonomy. This taxonomy can be used to begin categorizing large loads into like groups from the standpoint of potential grid stability challenges, but also ancillary service market participation. It is a first step toward mapping the operating characteristics of large loads to the requirements of various ancillary services, as well as risk factors for grid instability. It is not a comprehensive description of the operating characteristics of large loads or a set of requirements that loads must conform to. Rather, the taxonomy creates a framework for establishing market participation models and/or interconnection requirements for loads falling into these categories. It is a flexible framework that allows for “migration” from one type of large load to another based on changes in how a load is equipped, managed, and operated.

^d On-site power provides the capability to detach partially or fully from the grid yet still maintain operations and thereby provide a critical service to the grid at times of stress.

^e For example, data center owners and operators that house and serve multiple tenants are generally not in a position to make flexible power commitments on behalf of their tenants.

As described previously, this taxonomy considers large loads as those that are at or above 50 MW. The taxonomy assumes that loads of the same size have the same stability implications across different interconnections. It is also assumed that loads have the same considerations whether they are interconnecting at the transmission or distribution level. These assumptions may prove inaccurate in practice, however consideration of such finer details is left for future work.

The large load taxonomy is presented below in Figure 3. Based on the three fundamental characteristics of load size, load variability, and load flexibility, the taxonomy can result in 12 basic large load types. For example, a slow, inflexible load that has a peak demand of 200 MW would be a **Type B-2**, while a fast, flexible load with peak demand of 80 MW would be a **Type C-1**.

	Load Variability	Load Flexibility	Load Size		
			50-99 MW	100-999 MW	>1,000 MW
Type A	Slow	Flexible	1	2	3
Type B	Slow	Inflexible	1	2	3
Type C	Fast	Flexible	1	2	3
Type D	Fast	Inflexible	1	2	3

Figure 3. Large load taxonomy.

5. APPLYING THE TAXONOMY

Large loads of any type can be examined using this taxonomy to determine how they behave relative to other commercial electricity customers. Over the past 20 years, data centers have emerged as a very notable and unique type of load, but there are many other large industrial facility types that have long existed and, like data centers, are steadily increasing in size and potential impact on the grid. Balancing authorities, utilities, regulators and other stakeholders should have a firm understanding of important distinctions amongst these large loads, and this taxonomy can serve as a helpful starting point. Provided below is a brief discussion of some of these different large load types.

TRADITIONAL LARGE LOADS

While the proliferation of data centers is dominating many business, policy and regulatory discussions regarding large loads and the grid, there remains a very wide assortment of so-called traditional or legacy large loads that must also be considered. These typically include very large, energy intensive operations like oil refining, mining, and metals production. In addition, there are newer types of industrial facilities that increasingly utilize very large amounts of energy from the grid, such as hydrogen production plants that use electrolysis to separate hydrogen from water.

DATA CENTERS

The fundamental role of any data center is to provide large quantities of reliable data storage, management and/or processing, but these functions are applied in different ways and for different end results depending on the type of data center. For example, some data centers are dedicated to a single organization's use (enterprise data centers) while some exist to serve numerous clients (colocation data centers). Some are designed to be smaller (edge and modular data centers) and some seek to maximize size and computing power (hyperscale data centers). For purposes of applying the large load taxonomy, it can be useful to consider the different computing services being provided by data centers, noting that some data centers may provide multiple services or may switch between primary computational services performed. While not an exhaustive list, the following captures some of the main functions that larger data centers can provide^f:

- **Cloud Services:** These include the vast array of on-line services that businesses, governments, and consumers utilize daily, such as financial transactions, on-line sales, business record management, and entertainment. Most of the current data center industry is focused on this type of service, which generally requires constant and consistent computing and energy use.
- **Artificial Intelligence (AI) Training:** Data centers are increasingly needed to develop and refine the models that are used in a growing array of AI applications. This function requires large amounts of computing power that is typically applied in large batches with very fast fluctuations (i.e., high variability) in energy demand.
- **AI Inference:** Deployed AI models are utilized to generate results for many types of AI applications, and this function must also occur within data centers. Like training, AI inference requires significant computing power but typically does not involve the same large and rapid operational and power demand fluctuations.
- **Cryptocurrency Mining:** Like the other data center functions described above, cryptocurrency mining requires facilities with large, dedicated computing power used to solve complex calculations and facilitate financial transactions. Cryptocurrency mining is generally less sporadic from an operational and power consumption standpoint and can maintain a more flexible posture [2].

TAXONOMY EXAMPLES

The taxonomy can be applied to help distinguish these many diverse types of large loads and draw conclusions about a given facility's potential effects upon grid stability. Provided below are examples of representative large loads and where they are categorized within the taxonomy. As mentioned previously, the taxonomy is not intended to provide a rigid one-time classification, but to demonstrate where a large load may currently sit, based on available information, recognizing that recategorization is very possible with certain operational, technological and business changes.

Example 1: A hyperscale cloud services data center operates at a 400 MW capacity. The facility delivers continuous computing capacity for global services such as web browsing, streaming, and hosting enterprise software used for various business processes. Its baseline electricity demand is consistently high, resulting in a relatively predictable, static load profile that experiences very little fluctuation. Any decrease in external power automatically activates the data center's uninterruptible power supply (UPS) battery systems which can sustain the facility for several minutes, allowing for backup diesel-powered generators to cycle on and keep servers online for much longer periods of time without support from the grid. In the event of an impact on the grid, the facility communicates with the company's centralized,

^f There are various categorizations of data center types that focus on separate differentiators (e.g., design factors, size, function). NERC has often focused on the four-way functional split between cloud, AI training, AI inference, and cryptocurrency data centers [2].

integrated planning system which provides facility operators with hourly instructions to curtail non-urgent computing functions and reroute more important functions to one or several of the company’s data centers in other regions of the country. As a result, it provides grid operators with a stable base load that can still flex in response to contingencies.

This large load is a **Type A-2** within the taxonomy. Examples of actual facilities that may fit this profile include Google’s large data center in Council Bluffs, Iowa ([Council Bluffs, Iowa – Google Data Center Location](#)) and Amazon’s data center clusters in Virginia and Oregon ([Regional Data Center Clusters Power Amazon’s Cloud | Data Center Frontier](#)).

	Load Variability	Load Flexibility	Load Size		
			50-99 MW	100-999 MW	>1,000 MW
Type A	Slow	Flexible	1	2	3
Type B	Slow	Inflexible	1	2	3
Type C	Fast	Flexible	1	2	3
Type D	Fast	Inflexible	1	2	3

Figure 4. Large load type A-2 example.

Example 2: A data center owned and operated by the U.S. Government serves as a comprehensive data repository for the National Intelligence Community. It gathers, processes, stores, and transmits sensitive data on a global scale. The facility has a 65 MW peak load and runs on a continuous 24/7 schedule, with very little variation in electricity demand. Given the data center’s critical function as a primary location for receiving and processing national intelligence, it is highly limited in its flexibility to respond to changes in grid demand through curtailing its power consumption. Diesel-fueled generators provide emergency backup power which can sustain 100% of the data center’s operations for approximately three days in the event of a disruption to the grid. But this resilience is designed and intended only for emergency situations and not to be deployed in other circumstances.

This large load is a **Type B-1** within the taxonomy. An actual facility that may fit this profile is the U.S. National Security Agency’s Utah Data Center (<https://www.datacenterdynamics.com/en/news/us-national-security-agency-hybrid-cloud-environment-now-operational/>).

	Load Variability	Load Flexibility	Load Size		
			50-99 MW	100-999 MW	>1,000 MW
Type A	Slow	Flexible	1	2	3
Type B	Slow	Inflexible	1	2	3
Type C	Fast	Flexible	1	2	3
Type D	Fast	Inflexible	1	2	3

Figure 5. Large load type B-1 example.

Example 3: An aluminum smelter and alloy plant produces 200,000 tons of aluminum annually and has an average electricity demand of roughly 300 MW across its operations. The facility’s energy demand fluctuates widely across different stages of smelting and refining: from the intense bursts of electricity required during furnace ignition and melting, to lower but still significant consumption during refining, casting, and idle phases. These cycles create steep peaks and troughs in power usage, often in short

intervals, driving substantial variability in its load profile. The plant hosts a combination of on-site generators and a UPS system which together provide emergency backup power for its most critical digital control systems, but the sheer scale of energy needed for the smelting process makes providing full backup for the entire facility impractical. Unlike certain data centers or other more modular large loads, the aluminum plant has little operational flexibility; furnace cycles cannot be paused or ramped down without disrupting production, damaging equipment, or compromising material quality. The need for continuous thermal stability and metallurgical precision prevents operators from shifting or curtailing demand in response to external signals, locking the facility into a rigid electricity usage pattern despite its large and highly variable load swings.

This large load is a **Type D-2** within the taxonomy. Examples of actual facilities that may fit this profile include the Century Aluminum Plants in Mt. Holly, South Carolina ([Century Aluminum Company - Products & Plants - Mt. Holly, SC](#)) and Sebree, Kentucky ([Century Aluminum Company - Products & Plants - Sebree, KY](#)).

	Load Variability	Load Flexibility	Load Size		
			50-99 MW	100-999 MW	>1,000 MW
Type A	Slow	Flexible	1	2	3
Type B	Slow	Inflexible	1	2	3
Type C	Fast	Flexible	1	2	3
Type D	Fast	Inflexible	1	2	3

Figure 6. Large load type D-2 example.

Example 4: A large data center campus operates at a combined 200 MW capacity across its facilities, and provides network-neutral data hosting, colocation, and interconnection services to multiple tenants. Located near a major urban area, the campus also serves as a central hub for the exchange of internet traffic and hosts equipment for several major telecommunications operators. Designed to lease space, power, and connectivity to multiple tenants—ranging from enterprises to government agencies—the combined facilities maintain a steady power profile, as each tenant runs mission-critical workloads that require uninterrupted availability and guaranteed service levels. Campus facilities are equipped with UPS and generator capacity to meet continuous demand for their clients, but cannot accommodate large daily swings in power usage. Facilities are also highly limited in their ability to curtail or throttle tenant servers without violating contracts or service level agreements. As a result, the campus provides grid operators with a reliable, static demand footprint but offers minimal opportunity for demand-response participation or fast load adjustments.

This large load is a **Type B-2** within the taxonomy. Examples of actual facilities that may fit this profile include the Equinix DC Internet Exchange ([Equinix Ashburn - PeeringDB](#)) and other colocation hubs located on the Ashburn Data Center Campus ([Ashburn Data Centers - 150 Facilities from 3 Operators](#)).

	Load Variability	Load Flexibility	Load Size		
			50-99 MW	100-999 MW	>1,000 MW
Type A	Slow	Flexible	1	2	3
Type B	Slow	Inflexible	1	2	3
Type C	Fast	Flexible	1	2	3
Type D	Fast	Inflexible	1	2	3

Figure 7. Large load type B-2 example.

Example 5: A large hyperscale data center operates at a 1.5 gigawatt capacity, and houses technically advanced information technology (IT) infrastructure designed to handle compute-intensive tasks needed to train AI models. Unlike hyperscale cloud data centers (which run steady, customer-facing workloads), this AI training facility typically runs jobs in ‘batches’ which can cause large fluctuations in electric power demand. Training runs may last for days or weeks, but they also feature idle periods or ramp-down windows between jobs. At certain points in the training process, real power demand can fluctuate from anywhere between 10% and 70% of the data center’s peak capacity within the span of several minutes. Though it experiences large, rapid power fluctuations, the facility’s infrastructure—featuring containerized workloads, advanced job schedulers, and modular power distribution, also allows for flexible scheduling to manage AI jobs, including the ability to prioritize, queue, or delay certain jobs. Non-urgent tasks can be rapidly shifted in time to reduce peak load or align with low-cost/low-carbon energy availability. The result is a facility that couples periods of highly variable, compute-driven demand with an agile ability to shed or defer load.

This large load is a **Type C-3** within the taxonomy. Examples of actual facilities that may fit this profile include Meta’s planned AI Training Data Center in Richland Parish, Louisiana ([Meta's 10B New AI Data Center in Richland Parish, Louisiana](#)) and the Vantage Frontier Campus project in Shackelford County, Texas ([Vantage Data Centers: Shackelford Data Center](#)).

	Load Variability	Load Flexibility	Load Size		
			50-99 MW	100-999 MW	>1,000 MW
Type A	Slow	Flexible	1	2	3
Type B	Slow	Inflexible	1	2	3
Type C	Fast	Flexible	1	2	3
Type D	Fast	Inflexible	1	2	3

Figure 8. Large load type C-3 example.

Example 6: A 150 MW capacity cryptocurrency facility hosts tens of thousands of mining rigs—specialized computing systems that perform complex mathematical calculations that validate and secure digital currency transactions. These mining rigs are purpose-built and made to perform very specific, continuous calculations that in turn produce a near-constant, predictable baseline electricity demand on a day-to-day basis. Since the profitability of cryptocurrency mining is intrinsically linked to the cost of electricity, load flexibility is built into the facility’s design. During periods of high electricity demand, mining operators actively work to reduce power consumption in response to increased electricity prices. Mining rigs are deployed in a modular design, allowing the facility to quickly curtail or scale down operations within seconds to minutes by powering down or throttling select racks of miners without

jeopardizing data integrity. As a result, the facility provides both a steady, dependable base load and maintains the capacity to rapidly respond to changes in grid demand.

This large load is a **Type A-2** within the taxonomy. An actual facility that may fit this profile is the Helios Bitcoin Mining Facility in Dickens County, Texas ([Galaxy To Acquire Helios Bitcoin Mining Facility From Argo Blockchain | Galaxy](#)).

	Load Variability	Load Flexibility	Load Size		
			50-99 MW	100-999 MW	>1,000 MW
Type A	Slow	Flexible	1	2	3
Type B	Slow	Inflexible	1	2	3
Type C	Fast	Flexible	1	2	3
Type D	Fast	Inflexible	1	2	3

Figure 9. Large load type A-2 example.

6. CONCLUSION

The purpose of this work is to address the growing need for a common approach to characterize large loads within the U.S. power system. By developing a shared taxonomy, this work seeks to provide a foundation for utilities, grid operators, and other stakeholders to better understand and manage the challenges and opportunities posed by integrating novel large loads. This paper has identified the essential characteristics of large loads and proposed a basic taxonomy to categorize them based on these characteristics. The authors acknowledge that there are nuances within large load behaviors and impacts that are not captured with the simple classifications provided by this taxonomy. The goal was to provide basic building blocks for common understanding, with the expectation that future work will provide more characterization of factors such as voltage, size relative to local system, presence and type of behind-the-meter co-generation, load criticality, and operational visibility. Consideration of these factors will be needed to fully address NERC’s recommendations for large load characterization in a way that can be used for regulatory work.

This taxonomy will be used by the authors to serve as a framework within which the capabilities and requirements for different classifications of large loads to participate in grid support services can be explored. This includes participation through open-market structures and potentially as requirements of standardized interconnection agreements. By providing a structured approach, the taxonomy will help in identifying the specific services that different classes of large loads may be able to provide and identifying risks for those that cannot.

In the future, the classification system will be expanded so that more nuanced behaviors, such as backup power capacity and load criticality, can be captured within each classification. How these nuanced operating characteristics affect the provision of grid support services will also be considered. For load types that prove unsuitable for the provision of ancillary services, this framework will be used for the development of risk assessments and the development of risk mitigation strategies.

Load classification is not a static descriptor. Indeed, the essential characteristics of these loads will vary based on what equipment is installed at a given facility and how it is being used. If a load decides to enroll in a demand response program, adds new behind-the-meter generation or storage assets, adds additional capacity with new infrastructure, or changes some of its business functions (e.g. a data center starts processing more AI training requests), its classification and the corresponding interaction with the grid could change. In this way, this taxonomy also serves as a tool for large load owners and operators to develop facility expansion plans to target specific capabilities.

This taxonomy offers a foundational reference point for understanding the diverse behaviors of large loads and their implications for grid operations. By establishing a shared vocabulary, it enables more consistent communication across utilities, regulators, and asset owners, and supports more coordinated planning and policy development. As the grid continues to evolve, this framework can help stakeholders assess the capabilities and risks associated with emerging load types, guide infrastructure investment decisions, and inform the design of programs that align load behavior with system needs.

7. REFERENCES

1. Enriquez-Contreras, L. F., et al. 2025. “MegaWatt Mayhem: Grid Operator Challenges with Large Loads,” INL/RPT-25-87430. Idaho National Laboratory, Idaho Falls, ID.
2. NERC. 2025. “Characteristics and Risks of Emerging Large Loads: Large Loads Task Force White Paper.” North American Electric Reliability Corporation, Atlanta, GA.
https://www.nerc.com/comm/RSTCReviewItems/3_Doc_White%20Paper%20Characteristics%20and%20Risks%20of%20Emerging%20Large%20Loads.pdf.
3. MISO. 2023. “MISO Load Interconnection Whitepaper.” July 2023. Midcontinent Independent System Operator, Carmel, IN.
<https://cdn.misoenergy.org/MISO%20Load%20Interconnection%20Whitepaper629693.pdf>.
4. CPS Energy. n.d. “Large Load Customer Growth.” CPS Energy, San Antonio, TX. Accessed September 29, 2025. <https://www.cpsenergy.com/en/construction-and-renovation/latest-updates/large-load-customer-growth.html>.
5. Blevins, B. 2022. “Large Flexible Load Taskforce Interconnection and Planning working session.” Presentation, Electric Reliability Council of Texas, Austin, TX.
6. Miller, R. 2024. “The Gigawatt Data Center Campus is Coming.” Data Center Frontier, Accessed September 29, 2025. <https://www.datacenterfrontier.com/hyperscale/article/55021675/the-gigawatt-data-center-campus-is-coming>.
7. Kunder, P. 1994. “Power System Stability and Control”, McGraw-Hill.