

Review

Not peer-reviewed version

Energy Storage Systems for AI Data Centers: A Review of Technologies, Characteristics, and Applicability

[Saifur Rahman](#)* and [Tafsir Ahmed Khan](#)

Posted Date: 29 December 2025

doi: 10.20944/preprints202512.2564.v1

Keywords: AI data centers; energy storage systems; hybrid energy storage systems; high-variability electrical loads; GPU-intensive computing



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Review

Energy Storage Systems for AI Data Centers: A Review of Technologies, Characteristics, and Applicability

Saifur Rahman * and Tafsir Ahmed Khan

Advanced Research Institute, Virginia Polytechnic Institute & State University, Arlington, VA 22203, USA

* Correspondence: srahman@vt.edu

Abstract

The fastest growth in electricity demand in the industrialized world will likely come from the broad adoption of artificial intelligence (AI) - accelerated by the rise of generative AI models such as OpenAI's ChatGPT. The global "data center arms race" is driving up power demand and grid stress, which creates local and regional challenges because people in the area understand that the additional data center-related electricity demand is coming from faraway places, and they will have to support the additional infrastructure while not directly benefiting from it. So, there is an incentive for the data center operators to manage the fast and unpredictable power surges internally so that their loads appear like a constant baseload to the electricity grid. Such high-intensity and short-duration loads - which can have very high tariffs - can be served by Hybrid Energy Storage Systems (HESS) combining multiple types of storage technologies. Each storage type can respond to different timescales of load variations, enabling fast power balancing and reduced battery degradation. Properly designed HESS architectures, and their optimal operations, will make AI data center loads as baseloads, which will help the data center operators avoid high-tariff electricity. This review paper presents an overview of current energy storage technologies, their classifications, and recent performance data, focusing on their suitability for AI-driven computing. Technical requirements such as fast response, modularity, and long cycle life - that are critical for sustainable and reliable data center operations - are discussed. Special attention is given to Northern Virginia, the world's largest data center hub, where rapid load growth and grid limitations highlight the urgent need for hybrid storage deployment. The paper concludes by identifying key challenges, research gaps, and future directions for integrating hybrid energy storage with on-site generation to optimally manage the high variability in the data center load and build sustainable, low-carbon, and intelligent AI data centers.

Keywords: AI data centers; energy storage systems; hybrid energy storage systems; high-variability electrical loads; GPU-intensive computing

1. Introduction

1.1. Background and Motivation

The recent explosion of artificial intelligence (AI) applications has reshaped computing around the world. Each generation of AI model requires more computing power, faster movement of data, and higher reliability. The field has rapidly expanded, leading to the explosive growth of data centers globally. AI data centers, built to support large-scale model training and inference, have correspondingly consumed large quantities of electricity on par with small cities [1]. Generative AI has further accelerated this trend. State-of-the-art models like OpenAI's ChatGPT, Google Gemini, and Anthropic's Claude train on and run thousands of graphics processing units (GPUs) or tensor processing units (TPUs) in parallel [2]. In recent months, technology giants have entered a period of fierce competition that many analysts call an "AI arms race" [3]. Leading firms such as Google, Meta, Amazon, and OpenAI are investing hundreds of billions of dollars to increase computing power and

data center capacity with the goal of outcompeting rivals in the training and deployment of large AI models [4–6]. This arms race has accelerated the construction of new hyperscale facilities and is rapidly driving global electricity demand to levels not seen in recent decades, despite energy and climate concerns from environmental groups against building new data centers [7].

Due to this worldwide expansion, new data center clusters are appearing in many regions and competing for available power and infrastructure capacity [8–10]. Each cycle of AI training requires a considerable amount of energy and generates significant heat that needs to be continuously managed. To meet this increasing demand, cloud service providers and technology companies are rapidly constructing hyperscale campuses across North America, Europe, and Asia [11]. Among all regions, Northern Virginia is notable as the world's largest data center hub, hosting more than 250 facilities and several gigawatts of connected load [12,13]. Local utilities and communities are facing the challenge of supporting this massive demand while ensuring grid reliability [10]. Generation and transmission expansion, as well as required substation upgrades, take years to develop, but the need for AI is growing far more quickly. Similar grid bottlenecks are becoming serious problems for both operators and planners in other major hubs like Oregon (USA), Dublin (Ireland), and Singapore [14–16].

According to the International Energy Agency (IEA), data centers accounted for 415 TWh of electricity consumption in 2024, which is equivalent to about 1.5% of the world's total consumption [17]. About 45% of this usage came from the United States, with China coming in second at 25% and Europe at 15% [17]. A high-growth scenario would result in an annual electric power requirement of more than 219 GW by 2030 [18]. Various sources predict that the worldwide need for data center capacity might grow at a compound annual growth rate (CAGR) between 15% and 22% through 2030 [18].

In the US, data centers in 2024 consumed more than 4% of the country's electric energy [19,20]. This share is projected to rise to nearly 12% by 2030 [19,20]. To accommodate this rapid growth, an estimated 47 GW of additional generation capacity will be required by the end of the decade [21]. Meeting this power demand could require close to \$50 billion in new investments in US power generation infrastructure over the same period [21].

Generative AI workloads are expected to be a major contributor to this sharp rise in demand. Under high-growth projections, global data center capacity is expected to expand at a CAGR of about 22% by 2030. Over the same period, generative AI workloads are forecast to grow much faster, at roughly 39% CAGR, compared with around 16% CAGR for other digital services [18]. This gap between AI and other digital workloads highlights how AI-related computing is going to shape future power requirements. As AI-computing turns into a new form of industrial infrastructure, managing its energy footprint has become both a technical challenge and a social concern. Many communities recognize that while they host large data centers and support grid upgrades, most of their economic benefits often glow elsewhere. This creates a growing need for equitable and long-term energy management solutions.

1.2. Challenges of AI Data Center Power Profiles

AI data centers are significantly different from traditional data centers in terms of electricity. Traditional data centers that host cloud storage or web servers typically show reasonably stable loads throughout the day. In comparison, AI data centers operate with highly dynamic load profiles because of much higher rack power densities. Traditional server racks typically operate in the range of 4 kW to 12 kW, while high-density racks can go up to 20 kW [22]. In contrast, modern AI training racks generally draw between 50 kW and 200 kW, which highlights the enormous power requirements of GPU-based systems [22,23]. This power demand can fluctuate rapidly depending on whether training or inference tasks are running, resulting in significant variations in both electricity use and cooling demand. Such fast-changing loads may also introduce power quality issues such as harmonic distortion and voltage flicker [24]. To maintain the reliable operation of sensitive computing equipment, these power quality issues must be addressed.

The total power drawn by an AI data center can change dramatically within short time intervals [25]. During model training, thousands of GPUs may operate simultaneously, which pushes total consumption close to the facility's design limit for long durations. When the training jobs are finished, the load drops almost instantly [26]. Similarly, inference workloads generate short but intense surges in power demand. These rapid transitions can last only seconds or minutes but may reach several megawatts in magnitude, placing continuous stress on both the electrical and thermal management systems [25].

Such unpredictable fluctuations in load also create challenges beyond the facility boundary. For instance, utilities may experience serious voltage variations, transformer overloading, and frequency instability in their system due to the sudden changes in load. To manage these problems, utilities often maintain extra reserve capacity that raises operational costs and reduces efficiency. Like utilities, AI data center operators also experience an increase in operational cost due to the impact of peak demand charges. Most commercial electricity tariffs are based on the highest load in any 15 or 30-minute demand window, which means even brief spikes can lead to large increases in monthly bills [27]. As a result, AI data center operators are gradually seeking strategies to stabilize their loads and make their facilities appear as constant baseload customers to the grid.

As AI data centers become more power-hungry, traditional electrical design practices and backup systems are gradually becoming insufficient for modern AI workloads. Conventional UPS units and diesel generator-based backup systems were designed for rare outages, not for continuous cycling or multi-megawatt ramping throughout the day. This makes expansion planning and operations increasingly difficult for both the utilities and the operators. These challenges are not theoretical. They have already been observed in existing North American facilities, where real-time load ramps and oscillations have stressed both grid equipment and on-site power electronics. These observations emphasize the need for fast-responding energy storage that can buffer large swings in power demand and protect both the grid and internal systems. The following subsection presents examples that illustrate the magnitude and speed of these load variations in real AI data center operations.

1.3. Examples of Rapid and Oscillatory Load Patterns in Modern AI Data Centers

Recent field observations from North American facilities show how quickly large computing loads can change in practice [25]. Ref [25] shows an example where a data center ramps down from roughly 450 MW to about 40 MW within 36 seconds at a particular hour. The load then remains at approximately 7 MW for several hours before ramping back to 450 MW within only a few minutes. Such behavior demonstrates how AI-driven compute clusters can rapidly shed or restore enormous amounts of load, far faster than traditional industrial processes.

North American Electric Reliability Corporation (NERC) recently reported that AI data centers often exhibit periodic, repetitive, and sustained load patterns that can introduce forced oscillations at subsynchronous frequencies during compute cycles. This behavior is mostly driven by the rapid switching between "compute" and "communication" phases in large synchronized GPU workloads, which was also verified by researchers and engineers from Microsoft, OpenAI, and NVIDIA [26]. These oscillations may arise naturally from coordinated GPU activity or from unintended interactions among power-electronic equipment. In [25], NERC presented another example that shows a load profile of an AI data center during training for two minutes in which power fluctuations within the 0.6-1.0 p.u. range and forced oscillations were observed. Such fluctuations can create significant stress on both the utility grid and on-site power electronic-based equipment.

An additional example can be found in the NERC report, which shows a cryptocurrency mining facility in North America that ramped down 298 MW within 25 seconds after a control fault [25]. The load also exhibited subsynchronous oscillations with a peak-to-peak amplitude of roughly 25 MW. Although crypto mining differs from AI workloads, the fundamental behavior is similar. That means it has high digital loads that can change extremely fast and create power-system instability if left unmanaged.

These examples demonstrate why large AI facilities, especially in regions such as Northern Virginia in the United States, require fast-responding energy storage to buffer extreme ramps, control oscillations, and maintain power quality for both the grid and the data centers.

1.4. Role of Energy Storage Systems

Different energy storage systems (ESS) are emerging as alternatives or complements to traditional UPS solutions. Modern ESS can not only provide reliable backup power during outages, but also actively support power quality improvement, fast load-following, and peak shaving during normal operation [28]. Their fast response capabilities allow them to absorb or deliver power within milliseconds and make them well-suited to reduce power fluctuations and mitigate the disturbances caused by AI workloads. At the system level, maintaining a stable and predictable demand profile requires flexible, high-response, variable-duration energy solutions that combine reliability support with advanced load-balancing capabilities.

By storing energy during low-demand periods and releasing it when the load increases, storage systems can smooth out rapid power fluctuations and reduce peak demand in AI data centers. This helps both the operator and the local utility. The operator benefits from lower electricity costs and improved resilience, while the utility sees a more predictable and grid-friendly load. On top of that, a well-designed ESS can support other goals such as frequency control, voltage regulation, and greater use of on-site renewable energy if available.

The storage needs of AI data centers go far beyond conventional backup applications as their load fluctuations occur across multiple timescales- from milliseconds to hours. No single storage technology can effectively handle all of these variations. Hybrid energy storage systems (HESS) address this limitation by combining multiple storage technologies, each optimized for a different timescale and power range. For example, supercapacitors or flywheels can handle very fast transient spikes but last only for milliseconds to seconds [29]. Lithium-ion or sodium-sulfur batteries can manage fluctuations that span several minutes or hours [30]. For longer-term balancing, technologies like flow batteries or other storage can be added. By coordinating these as different layers, hybrid systems can deliver both fast response and large capacity while minimizing degradation of the battery components. These features of hybrid systems allow AI data centers to appear as steady and predictable loads to the grid, even though their internal power demand may change rapidly.

Hybrid systems also improve economic performance. It enables optimal use of each component according to its strengths and reduces replacement costs by extending lifetimes. Furthermore, hybrid systems can participate in multiple energy markets simultaneously, such as peak shaving, demand response, and ancillary services that generate additional revenue streams.

1.5. Scopes and Contributions of this Review

This review paper provides a comprehensive overview of available energy storage technologies and their applicability to AI data centers. It aims to bridge two growing research areas: advanced energy storage and high-performance computing infrastructure. While many previous reviews have focused on grid-scale storage, electric vehicles, or renewable integration, few have examined the specific requirements imposed by AI-driven computing environments where sharp increase or decrease in electrical power demands are common.

The paper begins with a structured overview of major energy storage classifications, including electrochemical, electrical, mechanical, and electromagnetic systems. For each category, the basic operating principles, key performance characteristics, and practical deployment considerations are discussed. Performance indicators such as response time, power capability, energy duration, efficiency, cycle life, and footprint are emphasized. These metrics are especially important for managing the highly dynamic nature of AI workloads.

Next, the review examines the role of hybrid energy storage systems (HESS) in addressing multi-timescale load fluctuations in AI data centers. Common hybrid architectures and technology pairings are introduced, highlighting how fast-response devices and high-energy systems can be coordinated

to improve power stability, reduce stress on batteries, and support grid-friendly operation. The discussion focuses on the specific load patterns seen in GPU-intensive training and inference workloads, rather than on general energy storage applications.

Finally, the paper highlights several key technical and research gaps related to energy storage deployment in AI data centers. These include the need for better AI-specific load modeling, improved coordination between fast and slow-response storage layers, and control strategies that remain robust, transparent, and practical for real-world operations. Stronger integration across IT workloads, power electronics, cooling systems, and grid interfaces is also required, along with greater attention to cyber-physical security risks.

By synthesizing current knowledge across storage technologies, hybrid configurations, and AI-specific operational needs, this review provides a clear and application-driven perspective on how energy storage can support the next generation of AI data centers. The insights are intended to assist researchers, system designers, and planners in understanding both the capabilities and limitations of existing storage solutions in this rapidly evolving domain.

2. Overview of Energy Storage Systems

Energy storage systems have played a key role in power systems for decades. The rapid growth of AI data centers is making them even more crucial for managing fast and unpredictable power demands. A wide range of technologies have been developed to store and release energy across different timescales, durations, and power ranges. These technologies support various operational goals such as peak shaving, frequency regulation, ride-through protection, power smoothing, and providing backup power. However, no single storage solution fits all applications. Some storage technologies are designed for direct deployment inside the data center, while others are better suited for utility-side or grid-level installations to support the large and growing energy supply needs of AI facilities. Each technology has unique advantages, limitations, and cost structures, which is why it is important to understand how they can best support high-density and highly variable AI workloads. This section provides an overview of major energy storage classifications, key performance metrics, and the characteristics that determine their suitability for data center environments.

2.1. Classification of Major Energy Storage Systems

Energy storage technologies can be classified in several ways. For example, we can classify them according to their storage medium, response time, application, and discharge duration. However, one of the most widely used approaches is to classify them based on the form of energy stored. Figure 1 illustrates a detailed classification of energy storage systems based on the type of energy stored.

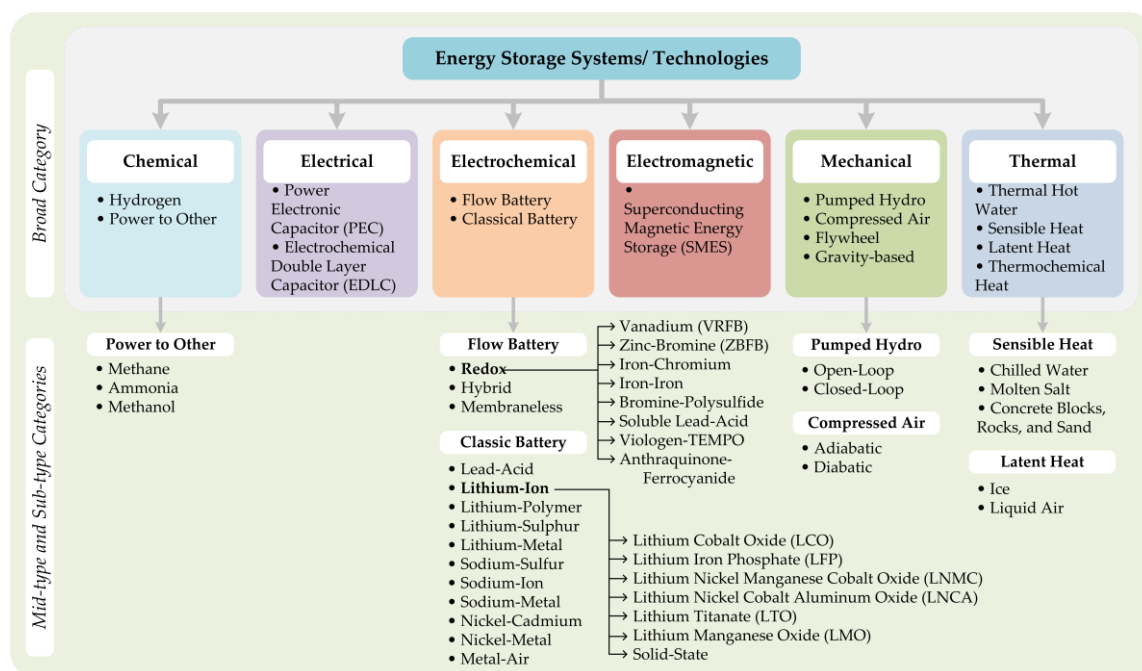


Figure 1. Classification of energy storage systems according to the form of energy stored.

2.2. Key Performance Metrics for AI Data Center Needs

Selecting suitable energy storage technology for an AI data center requires careful consideration of its characteristics, as these directly affect reliability, cost, and integration feasibility. The load pattern of AI data centers differs significantly from traditional data centers, and this makes certain energy storage parameters critical. The most important performance metrics for this application are summarized below.

- **Response Time:** AI workloads can change within milliseconds. ESS must respond instantly to maintain voltage stability and provide ride-through protection during short disturbances.
- **Power Rating:** GPU racks may draw 50-200 kW per rack. ESS should be capable of delivering high power output without performance loss.
- **Energy Density and Duration:** To support fluctuations that last from minutes to hours, ESS must have enough energy capacity to sustain these longer events.
- **Round-Trip Efficiency:** Frequent cycling for smoothing and peak shaving makes high efficiency crucial to keep energy losses and operating costs low.
- **Cycle Life and Degradation:** Frequent small cycles require storage technologies that can maintain long service life and stable performance, such as supercapacitors or flywheels.
- **Reliability and Safety:** Uptime requirements are strict in data centers. Also, ESS must meet fire safety standards, especially when placed inside buildings.
- **Space and Modularity:** Indoor deployments require compact and modular designs.
- **Cost and Value Stacking:** Technologies are assessed based on their capital cost and their ability to deliver multiple services, such as peak shaving, demand response, and power quality improvement.

Together, these metrics help AI data center operators and utilities to determine which storage technologies are best suited for dense and unpredictable loads, whether deployed inside data centers or on the grid side to support power delivery. It is obvious that no single technology can effectively meet all requirements. For this reason, hybrid storage configurations that combine multiple technologies are often needed. As illustrated in Figure 2, different energy storage technologies operate in different ranges of power and discharge duration. The left portion of Figure 2 is adapted from the “Energy Storage Technology Review” by SBC Energy Institute and updated using data from the DOE Global Energy Storage Database [31,32]. However, the right portion of the figure is created

entirely from the current database to reflect commercially available systems rather than theoretical values. This comparison emphasizes the importance of selecting and combining storage solutions that align with the unique load characteristics of AI data centers.

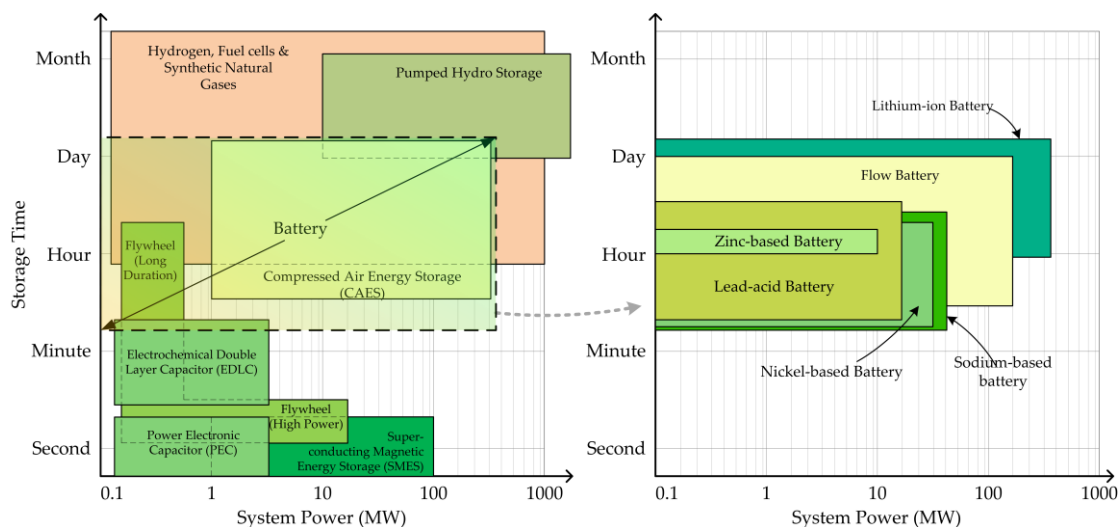


Figure 2. Energy storage technology map: storage time vs system power (left: all technologies, right: electrochemical ESS only).

Based on the requirements discussed above, electrochemical storage currently serves as the primary option for both on-site and grid-side support in AI data center applications. For this reason, it is an important category to examine first. The remaining subsections focus only on storage technologies that can be practically deployed on-site and combined in hybrid configurations to support fast, reliable AI data center operations.

2.3. Electrochemical Energy Storage

Electrochemical energy storage systems store energy in chemical form and release it through electrochemical reactions. They are the most widely deployed storage technology for data centers today due to their high round-trip efficiency, fast response times, modularity, scalability, and compact footprint suitable for indoor electrical rooms. Electrochemical storage technologies can be broadly categorized into classical batteries and flow batteries. Each category serves different operational needs and use cases.

2.3.1. Classical Batteries

Classical batteries such as lithium-ion, lead-acid, nickel-based, and zinc-based chemistries store energy within the cell structure. In most AI data centers, battery energy storage systems (BESS) are used as part of the uninterruptible power supply (UPS) [24]. They are installed both inside the facility and at the campus or electrical yard level, where compact size, fast response, and high reliability are important. Among all variants, lithium-ion (Li-ion) batteries have become the popular choice. They provide a good balance between high energy density, long cycle life, high efficiency, and low maintenance requirements. Their rapid millisecond-level response allows seamless ride-through protection and load balancing during GPU-based power surges. Several Li-ion chemistries are commercially available, each offering distinct trade-offs. For instance, lithium iron phosphate (LFP) batteries provide high thermal stability and long cycle life, which makes them well-suited for indoor UPS-type applications [33]. Lithium titanate oxide (LTO) offers exceptional cycling capability, extreme safety, and fast charging at the expense of lower energy density [33]. Nickel manganese cobalt (NMC) and nickel cobalt aluminum (NCA) variants achieve higher energy density but require stricter thermal management and safety controls, particularly in high-density AI environments [33].

On the other hand, conventional lead-acid batteries remain common in legacy UPS systems due to their low upfront cost and proven reliability for short-duration backup. However, they are less suitable for AI data centers that experience frequent and irregular load fluctuations due to their limited depth of discharge, shorter lifetime, and higher maintenance needs. Nickel-based and zinc-based batteries are used in niche applications but have seen limited adoption in modern AI-focused data centers.

2.3.2. Flow Batteries

Flow batteries store energy in liquid electrolytes housed in external tanks, while power conversion happens in the electrochemical cell stack [34]. This separation allows power (kW) and energy capacity (kWh) to be scaled independently [35]. More power can be provided by adding additional stack cells, while a longer duration is achieved by increasing the size of the electrolyte tanks [35]. As a result, 4-12+ hour discharge durations can be achieved through design choices rather than chemistry limitations [31,35]. Flow batteries typically offer long cycle life, low degradation, and high safety [35]. These characteristics make them well-suited for long-duration support in AI data center facilities.

Vanadium redox flow batteries (VRFB) are the most widely used type of flow battery. They avoid cross-contamination issues of common flow batteries and can maintain stable performance for 20-30 years or more of operation [34]. A schematic diagram of the working principles of a VRFB is shown in Figure 3. During operation, two liquid electrolytes containing different vanadium ions (V^{2+}/V^{3+} and V^{4+}/V^{5+}) are pumped through the cell stack and separated by an ion-exchange membrane. During charging and discharging, electrons move through an external circuit while ions pass across the membrane. This process allows energy to be stored and released repeatedly without degrading the active materials [34,36].




Figure 3. A schematic diagram illustrating the working principle of a Vanadium Redox-Flow Battery (VRFB).

Large-scale deployments are growing globally, especially in China, where multiple 100 MW-class systems are already in service [37]. Flow batteries have lower energy density and higher upfront cost than lithium-ion systems. However, they are a strong option for grid-side or campus-level storage that supports AI data center loads by enabling peak reduction, renewable integration, and improved resilience during extended GPU-intensive operations.

2.4. Electrical Energy Storage

Electrical energy storage technologies store and release energy directly in electrical form, which allows extremely fast response and high power delivery. They are especially useful for short-duration events that require rapid load compensation, such as the sudden GPU power ramps frequently observed in AI computing. In data center applications, electrical energy storage is mainly represented by Power Electronic Capacitors (PECs) and Electrochemical Double-Layer Capacitors (EDLCs), commonly known as supercapacitors (SCs).

2.4.1. Power Electronic Capacitors (PECs)

Power electronic capacitors, such as film capacitors and aluminum electrolytic capacitors, are widely used in power conversion systems. They provide extremely fast response times in the microsecond range ($\approx 0.1\text{-}100\ \mu\text{s}$) and can handle high power densities exceeding $10\ \text{kW/L}$ [31,32]. However, the amount of energy they can store is usually very small, often below a few watt-hours (Wh), even for large banks. This limits their role to power-conditioning tasks rather than true energy storage. PECs improve DC bus voltage stability, reduce harmonics and switching ripples, and help maintain power quality for GPU clusters and power electronic converters.

2.4.2. Electrochemical Double-Layer Capacitors (EDLCs)

Electrochemical double-layer capacitors (EDLCs), commonly known as supercapacitors (SC), respond within 1-10 milliseconds, which is much faster than lithium-ion batteries [38]. They achieve a very high power density of about $10,000\ \text{W/kg}$ and support extremely high cycling performance, typically from 500,000 up to more than 1,000,000 charge-discharge cycles with minimal capacity fade [39,40]. Their efficiency is also high, usually between 90-98% [38-41].

The operating principle of SC is the formation of an electric double layer. Charge is stored at the interface between the porous electrode and the electrolyte, where ions accumulate without chemical reaction [36]. This surface process occurs within the double-layer region, and the electrode material attracts ions from the electrolyte, as shown in Figure 4. Because of this mechanism, the active surface area of the electrode essentially determines the capacitance and performance of SCs.

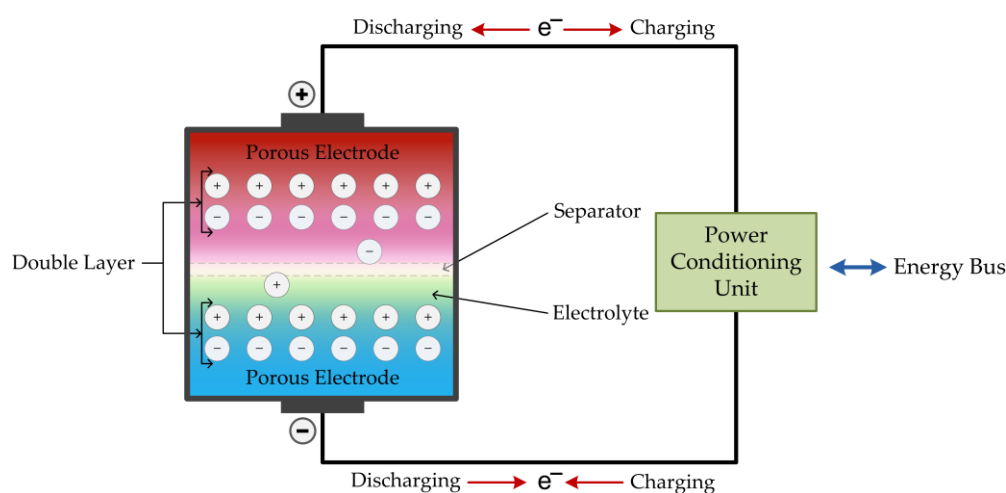


Figure 4. A schematic diagram illustrating the working principle of a supercapacitor.

The main limitation is energy density. Most supercapacitors provide only about 1-10 Wh/kg, compared with roughly 80-150 Wh/kg for lithium-ion batteries [39,40]. As a result, discharge durations are usually limited to seconds or a few minutes. This makes SCs ideal for high ramp-rate smoothing, short ride-through protection, and rapid response during AI workload spikes. Commercial modules can scale to the 1-2 MW range for grid applications. They provide enough

short-duration energy to protect battery systems from transient overloads and help stabilize power during sudden GPU surges.

PECs manage power-quality and voltage-stability issues at microsecond response speeds, while SC supercapacitors absorb rapid power bursts lasting milliseconds to seconds. Working together, they protect BESS from excessive cycling and help maintain stable power delivery for GPU-based AI infrastructure.

2.5. Mechanical ESS: Flywheel Energy Storage Systems (FESS)

Mechanical energy storage technologies convert electrical energy into potential or kinetic energy and can support both short-duration and long-duration power needs. They are widely used for grid support, renewable integration, and fast power balancing. While systems like pumped hydro (PHS) and compressed air (CAES) provide large-scale, long-duration storage, flywheel energy storage systems (FESS) have become increasingly relevant for high-power, short-duration applications inside or near AI data centers.

Flywheel energy storage systems store electrical energy as rotational kinetic energy in a high-speed rotor that is supported by magnetic bearings and housed in a vacuum chamber to reduce friction losses. As shown in Figure 5, the motor-generator interface and power electronic converter allow the rotor to accelerate during charging and decelerate during discharging. This mechanical process enables very fast response, typically within milliseconds. This makes flywheels highly effective for managing sudden GPU power spikes and stabilizing voltage during rapid AI workload changes.




Figure 5. Schematic diagram of a flywheel energy storage system.

FESS can supply power for a few seconds to a few minutes, which is long enough to ride through short disturbances or provide support until batteries or generators respond [42]. They offer high efficiency in the range of 90-95% and cycle life that often exceeds one million cycles, allowing them to handle frequent transients without degradation [36,43]. Although their energy capacity is smaller than battery or fuel cell ESS, flywheels act as a valuable fast-response buffer inside or near AI data centers. This reduces stress on primary high energy ESS, improves overall power quality, and supports continuous and reliable GPU computation.

2.6. Electromagnetic ESS: Superconducting Magnetic Energy Storage (SMES)

Superconducting magnetic energy storage (SMES) stores energy in the magnetic field of a superconducting coil that carries DC current with almost no electrical resistance. Energy is transferred through electromagnetic processes, giving SMES extremely fast response (microseconds), very high power output, and essentially unlimited charge-discharge cycle life [44,45]. A typical SMES

unit consists of a cryogenic coil, a DC-DC or DC-AC power conditioning system, and a bus interface, as shown in Figure 6 [46].




Figure 6. Generic structure of a superconducting magnetic energy storage (SMES) system.

These characteristics make SMES a strong high-power layer for hybrid systems in AI data centers. GPU clusters often create rapid, millisecond-scale load spikes, and SMES can absorb these disturbances without degrading over time. Although its energy capacity is limited, SMES significantly improves power quality and reduces stress on batteries, making it useful for stabilizing fast AI workload fluctuations.

3. Hybrid Energy Storage Systems (HESS)

AI data centers experience rapid and unpredictable changes in power demand due to GPU training and inference tasks. A single type of energy storage system cannot efficiently support these fast transients and longer-duration variations at the same time. Fast-response technologies such as supercapacitors and flywheels can handle short power bursts, but they cannot sustain energy delivery for minutes or hours. Batteries can supply energy for longer periods, but frequent high-power cycling accelerates their degradation. Although batteries are the most popular choice for data center applications, batteries alone are not enough to meet the demands of large-scale AI operations [47]. The next generation of data centers requires smarter energy systems, where storage technologies work together to deliver stability, resilience, and efficiency. Hybrid energy storage systems (HESS) combine multiple technologies so each one operates within its optimal performance range. This approach improves system efficiency and reduces battery wear. It also creates a smoother and more stable power profile for the grid, which makes HESS a practical and reliable solution for managing highly variable AI workloads in modern data centers.

3.1. HESS Architectures and Combinations

Hybrid energy storage systems can be arranged in different configurations depending on their operating goals and how the individual storage devices work together. There are different ways for the coupling of the energy storage in a HESS, and these choices determine how power is shared, how quickly each device responds to load changes, and how the system supports both on-site equipment and the utility grid [48,49].

For AI data centers, the structure of the HESS plays a key role in smoothing fast GPU power spikes, managing longer energy variations, and maintaining a stable and predictable demand profile. High-power storage (HPS) devices manage rapid fluctuations, while high-energy storage (HES) devices provide sustained support [50]. The overall architecture ensures that both operate together as a coordinated system. In addition to electrical storage, thermal energy storage (TES), such as chilled-water or ice-storage systems, can also be integrated as part of a broader hybrid solution. Because cooling accounts for about 38% of total data center energy use, TES can help shift or smooth

cooling demand. This allows TES to complement electrical storage and reduce overall peak load at the facility [51,52].

Figure 7 illustrates these ideas. Figure 7(a) presents a generic HESS architecture, showing how converters and the energy management system (EMS) coordinate multiple storage units. Figure 7(b) shows practical hybrid storage combinations suitable for AI data centers, where long-term energy devices such as BESS, flow batteries (FB), or fuel cells (FC) are paired with fast-response devices like supercapacitors, flywheels, SMES, or short-term BESS. These combinations allow hybrid systems to cover a wide range of power and energy needs, making them well-suited for the high variability of modern AI workloads.

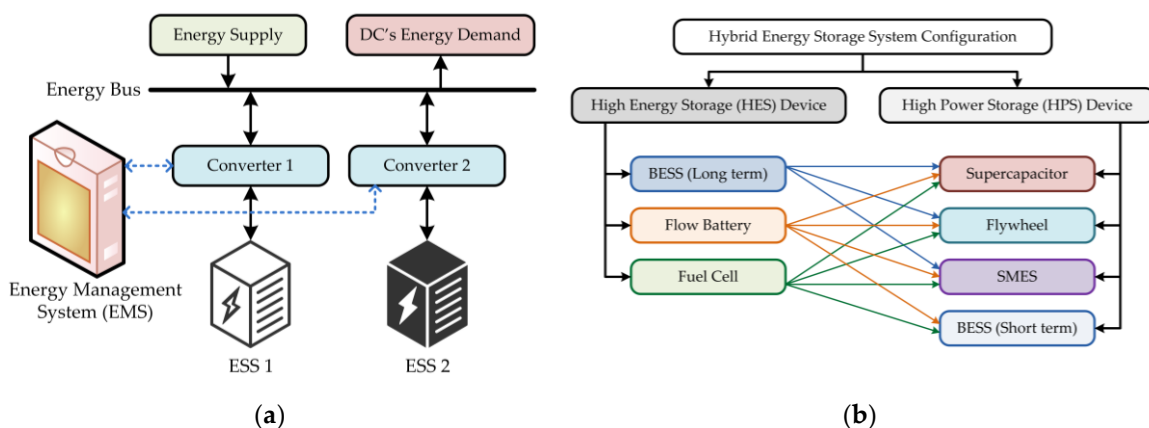


Figure 7. (a) Generic architecture of a hybrid energy storage system; (b) Practical HESS pairings tailored for AI data center needs.

3.2. Hybridization Benefits and Applications for AI Data Centers

AI data centers experience rapid, multi-timescale fluctuations in power demand driven by GPU training cycles, inference bursts, cooling transients, and software-scheduling patterns. A single storage technology struggles to manage this variability, especially when the power profile ranges from millisecond-level spikes to hour-long ramps. Hybrid energy storage systems (HESS) offer a practical solution by combining HPS technologies such as supercapacitors, flywheels, or SMES with HES systems such as BESS, FC, or FB. The HPS device absorbs fast disturbances, while the HES device supplies sustained energy. This division of roles allows the system to respond efficiently across all relevant AI timescales and reduces stress on the main storage layer. The key benefits of this hybridization in the context of AI data center operation are summarized below.

1. **Improved handling of fast and slow power fluctuations:** GPU clusters can transition from partial load to near-full consumption within seconds, and inference workloads often generate short but intense power bursts. HESS configurations assign these high-frequency events to fast-response HPS devices while reserving slower, multi-minute fluctuations for HES systems. This prevents over-cycling of HES and reduces the propagation of disturbances to the grid and internal power electronics.
2. **Extended battery lifetime and lower replacement costs:** Standalone battery systems face accelerated degradation in AI environments because of frequent and irregular workload-driven cycling [53]. Hybridization significantly improves battery life by shifting high-power, high-frequency demands to SC, flywheels, or SMES [40,42,54]. Studies in similar applications show that SC-BESS and SMES-BESS combinations can extend BESS lifetime by 19-26% [54,55]. A study shows that a flywheel-BESS hybrid configuration significantly slowed the battery aging by a factor of 300% [42]. This benefit alone can bring meaningful reductions in data center total cost of ownership.

3. **Power quality improvements for sensitive GPU loads:** Large GPU racks require stable voltage, low harmonic distortion, and a well-regulated DC bus. Fast storage layers in a HESS act as local “shock absorbers,” responding within microseconds to milliseconds to maintain voltage stability and dampen rapid load swings. This protects both IT equipment and upstream converters, which is critical in high-density AI racks operating at tens or hundreds of kilowatts.
4. **Accelerated interconnection and peak demand relief:** One of the largest barriers to new AI data center growth- especially in Northern Virginia in the United States- is the delay associated with securing firm grid interconnections [24]. HESS-supported battery strategies allow facilities to operate under interruptible interconnection agreements by riding through curtailments and temporary shortfalls. This approach can shorten interconnection timelines by years and unlock substantial new capacity.
5. **Lower energy costs and new revenue opportunities:** Hybrid systems lower operating costs by allowing batteries to perform energy arbitrage while high-power devices handle rapid fluctuations without cycling the main battery. This reduces peak demand charges and improves overall efficiency. Large HESS installations can also participate in wholesale market services such as frequency regulation, voltage support, spinning reserves, and fast frequency response, which can create additional revenue streams that help offset capital costs.
6. **Enhanced reliability and ride-through capability:** AI workloads cannot tolerate interruptions, as even brief disturbances can interrupt training jobs or damage equipment. HESS improves reliability by ensuring fast-response storage handles short-term disturbances while longer-duration storage maintains continuity during sustained grid events. In some designs, HESS can supplement or partially replace traditional UPS infrastructure [24].
7. **Scalable configurations for diverse data center needs:** Different AI campuses may prioritize peak shaving, interconnection acceleration, fast transient response, or long-duration energy shifting. HESS allows flexible pairing such as BESS-SC, BESS-FESS, BESS-SMES, FB-SC, or FC-BESS (short term), based on space constraints, cost, grid limitations, and workload characteristics. This adaptability positions HESS as a key component of future AI power architectures.

3.3. Control Strategies for Hybrid Energy Storage System

Control strategies play a central role in hybrid energy storage systems for AI data centers. Because AI workloads create rapid, multi-timescale power fluctuations, the controller (i.e., EMS) must decide how power is shared between HPS and HES devices. HPS units handle fast GPU surges, while HES units manage longer variations and overall energy balancing.

Most HESS controllers use a two-layer structure with a fast device-level controller that regulates converters and bus voltage, and a supervisory controller that assigns power based on frequency components, state-of-charge limits, or optimization rules [56]. Common methods include filtration-based power splitting, rule-based control, droop control, and deadbeat control [54–58]. Newer research applies MPC, fuzzy logic, neural networks, and optimization techniques to improve dynamic performance and storage lifespan [56,59–61].

For AI data centers, key control goals include smoothing millisecond GPU spikes, reducing battery cycling, improving power quality, and integrating with facility energy management. These needs require fast and precise coordination between storage devices. Although most existing methods were developed for microgrids or renewable systems, their core challenges, such as high-frequency power balancing and SOC protection, are still relevant. Table 1 summarizes representative control strategies from recent literature alongside the specific challenges they can address in AI data center environments.

Table 1. Brief overview of control strategies for HESS in applications that are relevant to data centers.

HESS	Control Method	Objective	Ref.
BESS-SMES	Modified droop control	Share power optimally, suppress voltage deviations, and protect the battery during rapid load change	[54]
BESS-SC	Filtration-based	Extending BESS lifetime	[55]
FB-SC	Filtration-based, Rule-based and Fuzzy logic-based	Handle fast load fluctuations while minimizing battery (FB) stress and reducing system cost	[57]
BESS-FESS	Moving-average filtering & fuzzy logic-based	Minimize battery ramp rate and preserves flywheel SoC	[58]
FC-short term BESS	Model predictive control	Reliably meet off-grid load while maintaining safe battery and hydrogen storage levels and maximizing long-term system efficiency over a 25-year horizon	[59]
BESS-TES	Optimization-based dispatch	Support the power grid and earn revenue for Data Center	[60]
BESS-TES	Optimization and model predictive control	Minimize operating cost, refining real-time charge-discharge operation, reducing SOC fluctuations and improving stability in Data Center	[61]
BESS-SC	Fuzzy Logic-based	Enhancing BESS performance during load changes	[62]
BESS-SC	Filtration-based	Preventing premature degradation of ESS	[63]

4. Discussions on Deployment Challenges and Research Opportunities

The previous sections showed that many energy storage technologies and hybrid architectures can support AI data centers. However, moving from technical potential to real deployments involves several practical challenges. These challenges relate to integration within the facility, coordination with the grid, and the need for better models and control strategies that are custom-made to handle AI workloads.

4.1. Integration and Design Challenges Inside the Data Center

At the facility level, one of the main challenges is how to integrate storage into existing power and cooling architectures. Indoor space in AI data centers is expensive and often prioritized for IT equipment. It is possible to house storage both inside and outside the building. In order to place large battery systems inside the building, careful planning for fire safety, ventilation, access routes, and compliance with local codes is required. When storage is installed outdoors, it must be coordinated with switchgear, transformers, and backup generators so that all systems operate safely as one unified plant.

Cooling is another integration constraint. High-density GPU racks already drive cooling systems close to their design limits. On top of that, poorly managed storage systems can add extra thermal loads and change the power consumption profile of chillers and pumps. Hybrid solutions that combine electrical storage with thermal energy storage can help, but they also increase the complexity of system design. Future work is needed on co-design methods that jointly optimize electrical and thermal storage for AI data centers rather than treating them as separate subsystems.

4.2. Grid Coordination, Siting, and Market Barriers

Rapid growth of AI data centers is putting increasing pressure on electric grids, especially in regions where transmission and interconnection capacity are already near their limits. In the United States, areas such as Northern Virginia clearly illustrate the tension between fast-paced AI expansion and limited transmission capacity [12,64,65]. New data center clusters often face long interconnection queues and may require accepting interruptible service or curtailment during peak periods. Energy

storage can help by absorbing or supplying power locally, which reduces stress on upstream grid equipment. However, this benefit depends on having clear rules that allow storage to participate in reliability services, demand response programs, and emergency operations.

Siting is also a key issue. In Northern Virginia, both existing and proposed data centers are facing growing community opposition because of water use for cooling, noise and low-frequency sound near residential areas, and broader environmental and land-use impacts [64,66]. Many locations are already constrained by zoning limits and limited developable land, which makes it difficult for data centers to allocate additional space for large storage systems. Because of the footprint, noise, or safety requirements, not all storage technologies are suitable for every site. These challenges reinforce the need for flexible siting strategies, including grid-side storage shared by multiple facilities, smaller modular systems, or hybrid solutions that reduce the on-campus footprint.

Another barrier is market and tariff design. AI data centers may fall under special customer classes or operate in restructured markets where capacity, energy, and ancillary services follow different rules and price signals [67,68]. This makes it difficult for AI data centers to deploy and operate storage systems and creates uncertainty about long-term revenue. Even hybrid storage systems that reduce load swings, support frequency regulation, or provide black-start capability may not be fully compensated for all the services they provide. Therefore, clearer policies are needed to recognize the multi-service value of storage and support grid-friendly operation of AI data centers.

4.3. Modeling, Control, and Cyber-Physical Research Gaps

Research on hybrid energy storage models and control strategies specifically tailored for AI data center applications is still limited. AI data center workloads follow software-driven patterns, have strong time correlation, and can change suddenly when training jobs start or stop. This calls for new models that link IT behavior, power electronics, cooling, and storage in a unified way.

Hybrid storage systems must coordinate fast millisecond-level devices with slower batteries and thermal systems. This makes simple rule-based control insufficient as systems grow larger and must support both on-site operations and grid services. Advanced methods such as model predictive control, reinforcement learning, and adaptive filtering are promising, but they must remain reliable and easy to understand for critical operations. Cyber-physical security adds another layer of challenge because failures or attacks on storage controls could cause large power swings that affect both the data center and the grid. Very little research has explored these AI-specific challenges so far, and there is a clear need for new studies that address modeling and robust control, and secure coordination across IT, storage, cooling, and grid domains.

4.4. Toward Practical Design Guidelines

Given these challenges, there is a clear need for practical design and planning guidelines. Such guidelines should support operators in addressing key questions, including:

- What combination of HPS and HES is most suitable for a specific AI workload profile?
- How should storage capacity be allocated between on-site systems and grid-side assets?
- Which control strategy offers an appropriate balance among performance and simplicity?

Developing such guidelines will require coordination among data center operators, utilities, regulators, and server providers. Relevant datasets should be made available to support future work, and early field experience with hybrid storage can help form standards, codes, and best practices that reduce risk and accelerate adoption.

5. Conclusions

AI data centers are emerging as one of the most demanding and dynamic electrical loads in modern power systems. Their rapid growth, high rack densities, and unpredictable multi-timescale load patterns place significant stress on both on-site infrastructure and the utility grid. These facilities can ramp up and down hundreds of megawatts within seconds and generate sustained oscillations

during GPU-intensive workloads. Conventional UPS systems - static diesel generator-based backup systems - and traditional planning and control methods are not well suited to manage this new class of load.

Energy storage systems, particularly when deployed in hybrid configurations, offer a practical and scalable pathway to address these challenges. Fast-response storage technologies such as supercapacitors, flywheels, and SMES can handle millisecond-level disturbances and protect sensitive power electronics. In contrast, high-energy systems such as lithium-ion, sodium-sulfur, flow batteries, and fuel cell ESS are used to manage longer-duration load changes and reduce peak demand. Hybrid energy storage systems integrate these complementary technologies into a coordinated platform that improves power quality, extends battery lifetime, supports accelerated interconnection, and enhances the reliability and resilience of AI campuses.

As AI data centers spread to new locations beyond established hubs like Northern Virginia in the United States, energy storage will become even more important for supporting safe, efficient, and grid-friendly growth. At the same time, the interaction between storage control, AI workload management, and power system planning opens many new research opportunities. These include real-time coordination across multiple system layers, secure operation of cyber-physical systems, integration with cooling systems, and improved frameworks for market participation.

Overall, the cases presented in this review show that hybrid storage architecture will be a foundational component of future AI energy systems. Their ability to handle both fast and slow power variations, support internal system stability, and meet external grid requirements will be essential as AI computing demands continue to grow over the coming decade. Hybrid systems also offer a natural path to integrate emerging on-site generation technologies. Continued research, real-world demonstrations, and coordinated policy development will be critical to turn these capabilities into reliable, cost-effective, and sustainable data center operations.

Author Contributions: Conceptualization, S.R. and T.A.K.; formal analysis, S.R. and T.A.K.; investigation, S.R. and T.A.K.; writing—original draft preparation, T.A.K.; writing—review and editing, S.R.; visualization, T.A.K.; supervision, S.R.. All authors have read and agreed to the published version of the manuscript.

Data Availability Statement: No new data were created or analyzed in this study.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Kimball, S. Data Centers Powering Artificial Intelligence Could Use More Electricity than Entire Cities. Available online: <https://www.cnbc.com/2024/11/23/data-centers-powering-ai-could-use-more-electricity-than-entire-cities.html> (accessed on 1 December 2025).
2. Li, T.; Pan, J.; Ma, S.; Aleksandr Raikov; Arkhipov, A. SimpleScale: Simplifying the Training of an LLM Model Using 1024 GPUs. *Applied Sciences* **2025**, *15*, 8265–8265, doi:<https://doi.org/10.3390/app15158265>.
3. Sigalos, M. OpenAI's Historic Week Has Redefined the AI Arms Race for Investors. Available online: <https://www.cnbc.com/2025/09/26/openai-big-week-ai-arms-race.html> (accessed on 1 December 2025).
4. Milmo, D. Boom or Bubble? Inside the \$3tn AI Datacentre Spending Spree. Available online: <https://www.theguardian.com/technology/2025/nov/02/global-datacentre-boom-investment-debt/> (accessed on 1 December 2025).
5. From OpenAI to Meta, Firms Channel Billions into AI Infrastructure as Demand Booms. *Reuters* 2025.
6. Barber, P. Data Centre Boom Sparks Deals Rush. Available online: <https://www.ft.com/content/42f3dec5-b8dc-49a2-aa5c-0e62ab529173/> (accessed on 12 December 2025).
7. Milman, O. More than 200 Environmental Groups Demand Halt to New US Datacenters. Available online: <https://www.theguardian.com/us-news/2025/dec/08/us-data-centers> (accessed on 12 December 2025).
8. Soni, A.; Sophia, D.M.; Navin, N. Microsoft Unveils \$23 Billion in New AI Investments with Big Focus on India. *Reuters* 2025.

9. Amazon Will Invest AU\$20 Billion in Data Center Infrastructure in Australia. Available online: <https://www.aboutamazon.com/news/aws/amazon-data-center-investment-in-australia> (accessed on 12 December 2025).
10. Chen, X.; Wang, X.; Colacelli, A.; Lee, M.; Xie, L. Electricity Demand and Grid Impacts of AI Data Centers: Challenges and Prospects. *arXiv (Cornell University)* **2025**, doi:<https://doi.org/10.48550/arxiv.2509.07218>.
11. Chapman, H. New Data Center Developments: December 2025. Available online: <https://www.datacenterknowledge.com/data-center-construction/new-data-center-developments-december-2025> (accessed on 13 December 2025).
12. JLARC. Data Centers in Virginia Available online: <https://jlarc.virginia.gov/landing-2024-data-centers-in-virginia.asp> (accessed on 12 December 2025).
13. Data Centers | Northern Virginia Regional Commission - Website Available online: <https://www.novaregion.org/1598/Data-Centers> (accessed on 12 December 2025).
14. Howland, E. Grid Constraints Limit Near-Term Data Center Growth in Northwest: NPCC Panelist. Available online: <https://www.utilitydive.com/news/data-center-load-northwest-npcc-power-plan-microsoft/735346/> (accessed on 12 December 2025).
15. Curran, I. New Data Centres Must Generate and Supply Electricity to Wider Market, Regulator Rules Available online: <https://www.irishtimes.com/business/2025/12/12/new-data-centres-must-generate-and-supply-electricity-to-wider-market-regulator-rules/> (accessed on 12 December 2025).
16. Er, D.; Ang, A. The Future of Data Centres in Singapore | Addleshaw Goddard LLP Available online: <https://www.addleshawgoddard.com/en/insights/insights-briefings/2025/real-estate/future-data-centres-singapore/> (accessed on 12 December 2025).
17. de-Bray, G.; Najeeb, N.; DeBlase, N. *AI and Energy Sectors More Intertwined than Ever.*; Deutsche Bank Research Institute, 2025;
18. Srivathsan, B.; Sorel, M.; Sachdeva, P. AI Power: Expanding Data Center Capacity to Meet Growing Demand Available online: <https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/ai-power-expanding-data-center-capacity-to-meet-growing-demand#/> (accessed on 12 December 2025).
19. Wendling, J. Charted: The Rising Share of U.S. Data Center Power Demand. Available online: <https://www.visualcapitalist.com/sp/gx03-charted-the-rising-share-of-u-s-data-center-power-demand/> (accessed on 12 December 2025).
20. U.S. Department of Energy. DOE Releases New Report Evaluating Increase in Electricity Demand from Data Centers. Available online: <https://www.energy.gov/articles/doe-releases-new-report-evaluating-increase-electricity-demand-data-centers> (accessed on 12 December 2025).
21. Davenport, C.; Singer, B.; Mehta, N.; Lee, B.; Mackay, J.; Modak, A.; Corbett, B.; Miller, J.; Hari, T.; Ritchie, J.; et al. *Generational Growth: AI, Data Centers and the Coming US Power Demand Surge.*; Goldman Sachs, 2024;
22. Johnstone, C. How Power Density Is Changing in Data Centers and What It Means for Liquid Cooling. Available online: <https://jetcool.com/post/how-power-density-is-changing-in-data-centers/> (accessed on 12 December 2025).
23. Bommarito, M. Rack Density Evolution: From 5kW to 350kW per Rack. Available online: <https://michaelbommarito.com/wiki/datacenters/technology/rack-density/> (accessed on 12 December 2025).
24. Taimela, P. Why Battery Energy Storage Is the Future of Data Center UPS Solutions. Available online: <https://www.flexgen.com/resources/blog/expert-qa-why-battery-energy-storage-future-data-center-ups-solutions> (accessed on 12 December 2025).
25. North American Electric Reliability Corporation (NERC). *Characteristics and Risks of Emerging Large Loads: Large Loads Task Force White Paper.* NERC, July 2025. Available online: https://www.nerc.com/globalassets/who-we-are/standing-committees/rstc/3_doc_white-paper-characteristics-and-risks-of-emerging-large-loads.pdf (accessed on 12 December 2025).
26. Choukse, E.; Warriar, B.; Heath, S.; Belmont, L.; Zhao, A.; Khan, H.A.; Harry, B.; Kappel, M.; Hewett, R.J.; Datta, K.; et al. Power Stabilization for AI Training Datacenters. *arXiv (Cornell University)* **2025**, doi:<https://doi.org/10.48550/arxiv.2508.14318>.

27. Rates and Tariffs | Virginia | Dominion Energy. Available online: <https://www.dominionenergy.com/virginia/rates-and-tariffs> (accessed on 12 December 2025).
28. Ali, Z.M.; Calasan, M.; Aleem, S.H.E.A.; Jurado, F.; Gandoman, F.H. Applications of Energy Storage Systems in Enhancing Energy Management and Access in Microgrids: A Review. *Energies* **2023**, *16*, 5930, doi:<https://doi.org/10.3390/en16165930>.
29. Aghmadi, A.; Mohammed, O.A. Energy Storage Systems: Technologies and High-Power Applications. *Batteries* **2024**, *10*, 141–141, doi:<https://doi.org/10.3390/batteries10040141>.
30. Liu, X.; Li, W.; Guo, X.; Su, B.; Guo, S.; Jing, Y.; Zhang, X. Advancements in Energy-Storage Technologies: A Review of Current Developments and Applications. *Sustainability* **2025**, *17*, 8316–8316, doi:<https://doi.org/10.3390/su17188316>.
31. Department of Energy Global Energy Storage Database. Available online: <https://gesdb.sandia.gov/> (accessed on 1 November 2025).
32. SBC Energy Institute. Building a Sustainable Energy System. Available online: https://wecanfigurethisout.org/ENERGY/Energy_home.htm (accessed on 1 November 2025).
33. Hannan, M.A.; Lipu, M.S.H.; Hussain, A.; Mohamed, A. A Review of Lithium-Ion Battery State of Charge Estimation and Management System in Electric Vehicle Applications: Challenges and Recommendations. *Renewable and Sustainable Energy Reviews* **2017**, *78*, 834–854, doi:<https://doi.org/10.1016/j.rser.2017.05.001>.
34. Dassisti, M.; Mastrorilli, P.; Rizzuti, A.; Cozzolino, G.; Chimienti, M.; Abdul-Ghani Olabi; Matera, F.; Carbone, A.; Ramadan, M. Vanadium: A Transition Metal for Sustainable Energy Storing in Redox Flow Batteries. *Elsevier eBooks* **2022**, 208–229, doi:<https://doi.org/10.1016/b978-0-12-815732-9.00128-5>.
35. Olabi, A.G.; Allam, M.A.; Abdelkareem, M.A.; Deepa, T.D.; Alami, A.H.; Abbas, Q.; Alkhalidi, A.; Sayed, E.T. Redox Flow Batteries: Recent Development in Main Components, Emerging Technologies, Diagnostic Techniques, Large-Scale Applications, and Challenges and Barriers. *Batteries* **2023**, *9*, 409, doi:<https://doi.org/10.3390/batteries9080409>.
36. Zhang, Z.; Ding, T.; Zhou, Q.; Sun, Y.; Qu, M.; Zeng, Z.; Ju, Y.; Li, L.; Wang, K.; Chi, F. A Review of Technologies and Applications on Versatile Energy Storage Systems. *Renewable and Sustainable Energy Reviews* **2021**, *148*, 111263, doi:<https://doi.org/10.1016/j.rser.2021.111263>.
37. Anyanwu, I.S.; Buzzi, F.; Peljo, P.; Bischi, A.; Bertei, A. System-Level Dynamic Model of Redox Flow Batteries (RFBs) for Energy Losses Analysis. *Energies* **2024**, *17*, 5324, doi:<https://doi.org/10.3390/en17215324>.
38. Liu, W.; Sun, X.; Yan, X.; Gao, Y.; Zhang, X.; Wang, K.; Ma, Y. Review of Energy Storage Capacitor Technology. *Batteries* **2024**, *10*, 271–271, doi:<https://doi.org/10.3390/batteries10080271>.
39. Gopi, C.V. V.M.; Ramesh, R. Review of Battery-Supercapacitor Hybrid Energy Storage Systems for Electric Vehicles. *Results in Engineering* **2024**, *24*, 103598, doi:<https://doi.org/10.1016/j.rineng.2024.103598>.
40. Yaseen, M.; Khattak, M.A.K.; Humayun, M.; Usman, M.; Shah, S.S.; Bibi, S.; Hasnain, B.S.U.; Ahmad, S.M.; Khan, A.; Shah, N.; et al. A Review of Supercapacitors: Materials Design, Modification, and Applications. *Energies* **2021**, *14*, 7779, doi:<https://doi.org/10.3390/en14227779>.
41. Luo, X.; Wang, J.; Dooner, M.; Clarke, J. Overview of Current Development in Electrical Energy Storage Technologies and the Application Potential in Power System Operation. *Applied Energy* **2015**, *137*, 511–536, doi:<https://doi.org/10.1016/j.apenergy.2014.09.081>.
42. Li, X.; Palazzolo, A. A Review of Flywheel Energy Storage Systems: State of the Art and Opportunities. *Journal of Energy Storage* **2022**, *46*, 103576, doi:<https://doi.org/10.1016/j.est.2021.103576>.
43. Zhang, J.W.; Wang, Y.H.; Liu, G.C.; Tian, G.Z. A Review of Control Strategies for Flywheel Energy Storage System and a Case Study with Matrix Converter. *Energy Reports* **2022**, *8*, 3948–3963, doi:<https://doi.org/10.1016/j.egy.2022.03.009>.
44. Hernando López de Toledo, C.; Munilla, J.; García-Tabarés, L.; Gil, C.; Ballarín, N.; Orea, J.; Iturbe, R.; López, B.; Ballarino, A. Design of Superconducting Magnetic Energy Storage (SMES) for Waterborne Applications. *IEEE Transactions on Applied Superconductivity* **2025**, *35*, 1–5, doi:<https://doi.org/10.1109/tasc.2025.3529414>.
45. Adetokun, B.B.; Oghorada, O.; Abubakar, S.J. Superconducting Magnetic Energy Storage Systems: Prospects and Challenges for Renewable Energy Applications. *Journal of Energy Storage* **2022**, *55*, 105663, doi:<https://doi.org/10.1016/j.est.2022.105663>.

46. Khaleel, M.; Yusupov, Z.; Nassar, Y.; El-khozondar, H.J.; Ahmed, A.; Alsharif, A. Technical Challenges and Optimization of Superconducting Magnetic Energy Storage in Electrical Power Systems. *e-Prime - Advances in Electrical Engineering Electronics and Energy* **2023**, *5*, 100223–100223, doi:<https://doi.org/10.1016/j.prime.2023.100223>.
47. BESS and Data Centers: Powering AI with Smart Energy Systems - CARRAR. Available online: <https://www.carrar.net/resources/bess-and-ai-driven-data-centers/> (accessed on 12 December 2025).
48. Atawi, I.E.; Al-Shetwi, A.Q.; Magableh, A.M.; Albalawi, O.H. Recent Advances in Hybrid Energy Storage System Integrated Renewable Power Generation: Configuration, Control, Applications, and Future Directions. *Batteries* **2023**, *9*, 29, doi:<https://doi.org/10.3390/batteries9010029>.
49. Bocklisch, T. Hybrid Energy Storage Approach for Renewable Energy Applications. *Journal of Energy Storage* **2016**, *8*, 311–319, doi:<https://doi.org/10.1016/j.est.2016.01.004>.
50. Hajiaghahi, S.; Salemnia, A.; Hamzeh, M. Hybrid Energy Storage System for Microgrids Applications: A Review. *Journal of Energy Storage* **2019**, *21*, 543–570, doi:<https://doi.org/10.1016/j.est.2018.12.017>.
51. Ahmed, K.M.U.; Bollen, M.H.J.; Alvarez, M. A Review of Data Centers Energy Consumption and Reliability Modeling. *IEEE Access* **2021**, *9*, 152536–152563, doi:<https://doi.org/10.1109/ACCESS.2021.3125092>.
52. Ahmed, K.M.U.; Alvarez, M.; Bollen, M.H.J. Reliability Analysis of Internal Power Supply Architecture of Data Centers in Terms of Power Losses. *Electric Power Systems Research* **2021**, *193*, 107025, doi:<https://doi.org/10.1016/j.epsr.2021.107025>.
53. Xu, B.; Oudalov, A.; Ulbig, A.; Andersson, G.; Kirschen, D.S. Modeling of Lithium-Ion Battery Degradation for Cell Life Assessment. *IEEE Transactions on Smart Grid* **2018**, *9*, 1131–1140, doi:<https://doi.org/10.1109/tsg.2016.2578950>.
54. Li, J.; Yang, Q.; Robinson, Francis.; Liang, F.; Zhang, M.; Yuan, W. Design and Test of a New Droop Control Algorithm for a SMES/Battery Hybrid Energy Storage System. *Energy* **2017**, *118*, 1110–1122, doi:<https://doi.org/10.1016/j.energy.2016.10.130>.
55. Gee, A.M.; Robinson, F.V.P.; Dunn, R.W. Analysis of Battery Lifetime Extension in a Small-Scale Wind-Energy System Using Supercapacitors. *IEEE Transactions on Energy Conversion* **2013**, *28*, 24–33, doi:<https://doi.org/10.1109/tec.2012.2228195>.
56. Babu, T.S.; Vasudevan, K.R.; Ramchandaramurthy, V.K.; Sani, S.B.; Chemud, S.; Lajim, R.M. A Comprehensive Review of Hybrid Energy Storage Systems: Converter Topologies, Control Strategies and Future Prospects. *IEEE Access* **2020**, *8*, 148702–148721, doi:<https://doi.org/10.1109/access.2020.3015919>.
57. Muhammad Hamza Ali; Darío Slaifstein; Federico Martin Ibanez; Zugschwert, C.; Pugach, M. Power Management Strategies for Vanadium Redox Flow Battery and Supercapacitors in Hybrid Energy Storage Systems. *2022 IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT-Europe)* **2022**, doi:<https://doi.org/10.1109/isgt-europe54678.2022.9960352>.
58. Maroufi, S.M.; Karrari, S.; Rajashekaraiah, K.; De Carne, G. Power Management of Hybrid Flywheel-Battery Energy Storage Systems Considering the State of Charge and Power Ramp Rate. *IEEE Transactions on Power Electronics* **2025**, *40*, 9944–9956, doi:<https://doi.org/10.1109/tpel.2025.3546013>.
59. Torreglosa, J.P.; García, P.; Fernández, L.M.; Jurado, F. Energy Dispatching Based on Predictive Controller of an Off-Grid Wind Turbine/Photovoltaic/Hydrogen/Battery Hybrid System. *Renewable Energy* **2015**, *74*, 326–336, doi:<https://doi.org/10.1016/j.renene.2014.08.010>.
60. Zhang, Y.; Tang, H.; Li, H.; Wang, S. Unlocking the Flexibilities of Data Centers for Smart Grid Services: Optimal Dispatch and Design of Energy Storage Systems under Progressive Loading. *Energy* **2025**, *316*, 134511–134511, doi:<https://doi.org/10.1016/j.energy.2025.134511>.
61. Wang, Z.; Yin, Z.; Yang, J.; Wang, J. Coordinated Optimization of Distributed Energy System and Storage-Enhanced Uninterruptible Power Supply in Data Center: A Three-Level Optimization Framework with Model Predictive Control. *Energy Conversion and Management* **2025**, *342*, 120137, doi:<https://doi.org/10.1016/j.enconman.2025.120137>.
62. Shayeghi, H.; Monfaredi, F.; Dejamkhooy, A.; Shafie-khah, M.; Catalão, J.P.S. Assessing Hybrid Supercapacitor-Battery Energy Storage for Active Power Management in a Wind-Diesel System. *International Journal of Electrical Power & Energy Systems* **2021**, *125*, 106391, doi:<https://doi.org/10.1016/j.ijepes.2020.106391>.

63. Ramos, G.A.; Costa-Castelló, R. Energy Management Strategies for Hybrid Energy Storage Systems Based on Filter Control: Analysis and Comparison. *Electronics* **2022**, *11*, 1631, doi:<https://doi.org/10.3390/electronics11101631>.
64. Patel, K.; Steinberger, K.; Debenedictis, A.; Wu, M.; Blair, J.; Picciano, P.; Oporto, P.; Li, R.; Mahoney, B.; Solfest, A.; et al. *Virginia Data Center Study: Electric Infrastructure and Customer Rate Impacts.*; 2024;
65. Blume, P. Dateline Ashburn: Data Centers Drive New Energy Disputes in Northern Virginia. Available online: <https://broadbandbreakfast.com/dateline-ashburn-data-centers-drive-new-energy-disputes-in-northern-virginia/> (accessed on 12 December 2025).
66. \$64 Billion of Data Center Projects Have Been Blocked or Delayed amid Local Opposition. Available online: <https://www.datacenterwatch.org/report/> (accessed on 12 December 2025).
67. Virginia State Corporation Commission. *Application of Virginia Electric and Power Company for a 2025 Biennial Review of Rates, Terms, and Conditions for Electric Service (Case No. PUR-2025-00058)*, 24 April 2025. Available online: <https://www.scc.virginia.gov/docketsearch/DOCS/84s201!.PDF> (accessed on 12 December 2025).
68. PJM Interconnection, L.L.C. *Ancillary Services Fact Sheet*. PJM Interconnection, Audubon, PA, USA. Available online: <https://www.pjm.com/-/media/DotCom/about-pjm/newsroom/fact-sheets/ancillary-services-fact-sheet.pdf> (accessed on 12 December 2025).

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.