Advances in Applied Energy 23 (2026) 100291

Contents lists available at ScienceDirect

Advances in Applied Energy

journal homepage: www.elsevier.com/locate/adapen

Grid integration of AI data centers: A critical review of energy storage solutions

Sina Mohammadi [a], Wayne Wang [b], Marcus Chen I Wada [a], Rouzbeh Haghighi [a], Ali Hassan [a], Hualong Liu [a], Archit Bhatnagar [b], Ang Chen [b], Wencong Su [a,*]

[a] University of Michigan-Dearborn, United States of America
[b] University of Michigan-Ann Arbor, United States of America

ARTICLE INFO

Keywords:
AI data center
Battery backup unit (BBU)
Battery energy storage system (BESS)
Fuel cell (FC)
GPU
Grid-interactive UPS (GiUPS)
Second-life battery (SLBESS)
Thermal energy storage (TES)

ABSTRACT

Artificial intelligence (AI) is driving unprecedented growth in data center (DC) scale and power demand. AI workloads impose highly dynamic, difficult-to-forecast power profiles on the utility grid, creating reliability and stability challenges that conventional DC architectures are not designed to address. This paper provides a critical review of energy storage systems (ESSs) as the key enabling technology for reliable grid integration of AI DCs. We organize the review around a four-layer hierarchical taxonomy, namely chip-level buffering, rack/server-level ESSs, facility-level uninterruptible power supply (UPS) systems, and grid-scale battery energy storage systems (BESSs), supplemented by non-battery technologies including fuel cells (FCs) and thermal energy storage (TES). Each layer is analyzed with respect to response timescale, power and energy ratings, operational role, integration challenges, and coordination requirements. Grid-interactive UPS (GiUPS) systems are examined in depth as a critical evolution of passive backup infrastructure into active grid-support assets capable of frequency regulation, fast frequency response, and voltage ride-through. Grid-scale BESSs are reviewed for their roles in peak shaving, renewable integration, and ancillary service provision. Second-life battery energy storage systems (SLBESSs) are evaluated as a cost-effective alternative for large-scale deployment. Key findings include: (i) AI DC load profiles differ fundamentally from traditional loads in their sub-second variability, making conventional ESS dispatch strategies insufficient; (ii) hierarchical, coordinated ESS deployment across all layers is necessary for effective load smoothing and grid support; and (iii) significant gaps remain in simulation tools, degradation modeling, load forecasting, and optimal multi-layer sizing. This review identifies open research challenges and future directions at the intersection of AI computing infrastructure and power system integration.

1. Introduction

Artificial Intelligence (AI) has fundamentally altered society's relationship with technology, spanning applications from simple search tasks to advanced scientific discovery. Large language models (LLMs) have been at the forefront of these high-complexity AI systems, with models such as ChatGPT and Claude becoming increasingly embedded in everyday use across the United States. These large-scale models are trained on trillions of text tokens collected from online sources and are deployed through web-based platforms that can attract millions of users over short periods of time. This surge in demand has prompted AI industry stakeholders to reconsider contemporary data center (DC) designs so that they can support the intensive computational requirements of training, inference, and fine-tuning for increasingly complex models. Energy requirements for AI queries are now nearly an order of magnitude greater than those of Google search, while AI training workloads have been observed to double approximately every 3.4 months [1]. Collectively, these trends highlight the need to rethink data center architectures to sustainably support the rapid scaling of AI workloads.

DCs are dedicated facilities, or collections of facilities, designed to house computational systems, telecommunications equipment, and data storage infrastructure. Energy demand within a DC is unevenly distributed: IT equipment and cooling systems account for the majority of electricity consumption, while auxiliary loads such as lighting and security represent a comparatively small fraction [2]. Traditional data centers (TDCs) have historically operated at total power levels typically below 30 MW [2], with individual electrical distribution branches commonly limited to a few tens of kilowatts. These facilities are predominantly connected to power distribution networks in

* Corresponding author.
E-mail address: wencong@umich.edu (W. Su).

https://doi.org/10.1016/j.adapen.2026.100291
Received 5 May 2026; Received in revised form 1 July 2026; Accepted 3 July 2026
Available online 16 July 2026

Nomenclature

Abbreviations

AI	Artificial Intelligence	LV	Low Voltage
AWS	Amazon Web Services	LDESSs	Long Duration Energy Storage Systems
BBU	Battery Backup Unit	LVRT	Low Voltage Ride Through
BESS	Battery Energy Storage System	MV	Medium Voltage
BMS	Battery Management System	MVDC	Medium Voltage Direct Current
BTM	Behind-the-Meter	OCP	Open Compute Project
CapEx	Capital expenditures	ORV3	Open Rack V3
CPU	Central Processing Unit	PCC	Point of Common Coupling
DC	Data Center	PDU	Power Distribution Unit
DMA	Direct Memory Access	PEMFCs	Polymer Electrolyte Membrane Fuel Cells
DoD	Depth of Discharge	PLL	Phase Locked Loop
DVFS	Dynamic Voltage and Frequency Scaling	PPA	Power Purchase Agreement
EMS	Energy Management System	PSA	Power Sharing Unit
ESS	Energy Storage System	RDMA	Remote Direct Memory Access
EV	Electric Vehicle	RES	Renewable Energy Sources
FC	Fuel Cell	RoCoF	Rate of Change of Frequency
FFR	Fast Frequency Response	RUL	Remaining Useful Life
FTM	Front-the-Meter	SEI	Solid Electrolyte Interphase
GFL	Grid-Following	SLB	Second Life Battery
GFM	Grid-Forming	SLBESS	Second-Life Energy Storage System
GiUPS	Grid-interactive Uninterruptible Power Supply	SoC	State of Charge
GPU	Graphics Processing Unit	SoH	State of Health
HBM	High Bandwidth Memory	SOFCs	Solid Oxide Fuel Cells
HDD	Hard Disk Drive	SSD	Solid State Disk
HPC	High Performance Computing	SST	Solid State Transformer
HVRT	High Voltage Ride Through	TDCs	Traditional Data Centers
IEC	International Electrotechnical Commission	TPU	Tensor Processing Unit
LFP	Lithium Iron Phosphate	TUPS	Traditional UPS
LLM	Large Language Model	VRT	Voltage Ride Through

Key Symbols

$C_{ch,t}$	Electricity price during charging at time step t (\$/kWh)		
$C_{ch/dis,t}$	Electricity price during charging/discharging at time step t (\$/kWh)		
C_{deg}	Degradation cost coefficient (\$/cycle)	t	Discrete time step (hours)
$P_{ch/dis}$	Charging/discharging power (kW)	Δf	Frequency deviation (Hz)

densely populated areas and operate under an architectural paradigm that emphasizes redundancy and exceptionally high uptime, making uninterrupted power delivery from the utility grid essential. TDCs are often sited near population centers to minimize latency for end users.

In contrast, AI DCs are driven by compute-intensive workloads, the widespread deployment of power-hungry GPUs, and the use of high-efficiency cooling technologies, all of which substantially increase power requirements. As a result, large-scale AI training facilities are increasingly located in remote or semi-rural regions where land, power, and cooling resources are more accessible and cost-effective. Rack-level power demand in AI DCs can exceed 100 kW, while aggregate facility consumption may scale to hundreds of megawatts or even gigawatt levels [3]. Consequently, large-scale AI training facilities often require direct interconnection with transmission networks and face stricter reliability, siting, and environmental constraints than traditional facilities [4]. Inference-oriented facilities, while potentially smaller in individual footprint, may connect at either the distribution or transmission level depending on aggregate scale and siting.

Beyond differences in scale and infrastructure, AI DCs also exhibit two distinct load profiles, training and inference, that differ fundamentally from those of TDCs. AI training workloads are characterized by rapid power fluctuations arising from idle periods, peak utilization events, and checkpointing operations, with significant load variations occurring on sub-second timescales [5,6]. In addition, training jobs often produce high-frequency, jitter-like power spikes. In contrast, inference workloads generally lack such rapid transients and exhibit comparatively smoother demand profiles. TDC loads, by comparison, behave more like conventional grid loads and are typically predictable using historical data and standard forecasting methods.

These distinctive load characteristics have important implications for grid operation. AI DCs are connected to the power grid through high-density power converters, which can increase power distortion at the utility level [7,8]. In addition, jittery spikes in training jobs may cause sub-frequency mechanical oscillations in synchronous generators, thereby increasing mechanical stress and accelerating corrosion. As a result, harmonic compensation and load smoothing become critical requirements for AI DC operation, both to improve load predictability and to mitigate adverse power-quality impacts on the power system.

Large-scale AI DCs also require the power grid to accommodate higher reserve margins and more proactive resource planning [9]. This is because forecasting their load dynamics is more challenging than for traditional loads, and the uncertainty associated with these facilities is significantly higher. Therefore, additional reserve capacity is needed both to maintain normal grid operation and to support grid expansion when integrating AI DCs.

Moreover, connecting or disconnecting large-scale AI DCs may jeopardize system stability because of high-range frequency fluctuations. These facilities are designed to rapidly transfer to backup power in response to even minor grid disturbances, which can trigger large and sudden load drops at the transmission level. Such disturbances may activate load-frequency relays or force grid operators to implement load-shedding actions. A recent transmission fault in Virginia triggered an abrupt 1500 MW reduction in system load as multiple DCs transitioned to backup power, leading to significant frequency and voltage

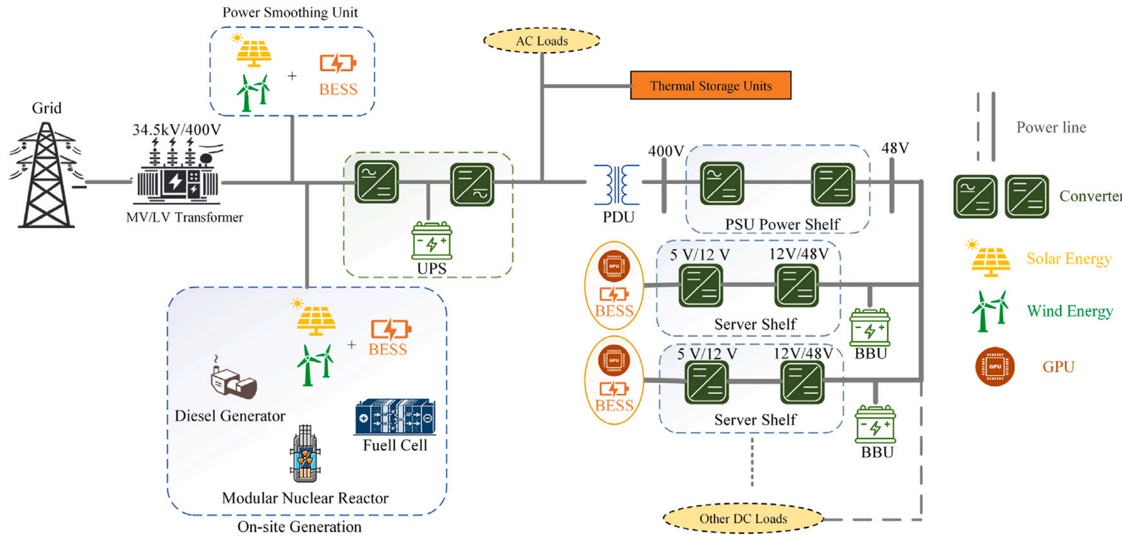

Fig. 1. Overall architecture of an AI data center with hierarchical energy storage systems. The figure illustrates the power delivery path from the utility grid through the medium-voltage transformer, UPS/GiUPS, power distribution units, and rack busbars to the IT equipment (GPUs/TPUs).

deviations that required operator intervention [10]. Events such as this illustrate the limitations of conventional grid planning and operational strategies when accommodating large AI DC loads.

The above challenges show why ESSs are especially important for AI DCs compared with TDCs. AI DCs require not only backup energy for reliability, but also fast and coordinated storage support to manage rapid training-load fluctuations, reduce upstream power-quality impacts, limit sudden grid-side ramping during backup transitions, and support reserve and frequency requirements at the transmission or distribution level. Therefore, ESSs can address multiple AI DC integration challenges simultaneously, including power smoothing, ride-through support, peak reduction, renewable-energy coordination, cooling-load shifting, and grid-interactive services. In this context, ESSs can be deployed across multiple physical layers and response timescales, ranging from GPU- or chip-level buffers for local power smoothing, rack-level battery backup units (BBUs), and UPS-embedded storage, to facility-level battery energy storage systems (BESSs) for data-center-level smoothing, co-generation, and grid services such as frequency regulation and peak shaving. Beyond battery-based solutions, thermal energy storage (TES) can support cooling flexibility, while fuel cells (FCs) have the potential to replace diesel backup generation as the technology matures. To the best of our knowledge, there is no comprehensive document specifically focused on these multi-layer support roles of ESSs for the reliable integration of AI DCs with the power grid. This broader role motivates a systematic review of ESS architectures across different physical layers, operational functions, and response timescales.

To provide a structured perspective on the diverse storage technologies reviewed in this paper, Table 1 summarizes the hierarchical ESS classification framework organized by storage layer, typical response timescale, power and energy ratings, primary application role, and the section in which each technology is discussed. This classification is intended to guide readers in comparing technologies across layers and to serve as a roadmap for the remainder of the paper. A more detailed cross-technology comparison, including cost and suitability metrics, is provided in Table 5.

Unlike prior reviews that separately examine DC energy demand, grid impacts, or ESS technologies, this paper provides an integrated ESS-focused review for AI DCs. In contrast to [9,11], our focus is not limited to characterizing the grid impacts of large AI DC loads, but rather on how ESSs can enable their reliable integration with the power grid. Moreover, unlike studies centered on ESS chemistry technologies [12,13], this review emphasizes storage architectures, operational roles, and coordination strategies across chip-, rack-, facility-, and grid-level layers. Specifically, the main contributions include a hierarchical ESS taxonomy for AI DCs, a grid-integration-oriented analysis of GiUPS and BESS coordination, and a cross-layer framework for understanding how different ESS technologies can jointly support AI DC reliability, load smoothing, and grid services. Therefore, this study provides a useful reference for academic and industry stakeholders seeking to explore the role of ESSs in future AI DCs.

2. AI DC power architecture

The AI DC infrastructure comprises different sections, including IT equipment, power distribution networks, cooling systems, and miscellaneous units such as lighting and security. In this section, we provide a brief introduction to the AI DC components and the ESS units in its structure. The overall structure of the AI DC with all ESSs is provided in Fig. 1. At a high level, ESSs in AI DCs can be categorized according to their physical location and operational timescale. Chip- and server-level storage mainly addresses fast local power transients, rack-level BBUs provide distributed ride-through and fault containment, facility-level UPS/GiUPS systems protect critical loads and enable grid-interactive operation, grid-scale BESSs support power smoothing and longer-duration grid services, while TES and FCs address cooling flexibility and low-carbon backup generation, respectively. This layered view is used throughout the paper to avoid treating ESS technologies as isolated solutions and to emphasize their complementary roles in AI DC grid integration.

2.1. IT equipment

The IT equipment itself comprises high-performance computational electronic chips such as GPUs, TPUs, and CPUs. It also includes advanced storage architectures, such as solid-state storage units, for storing required information with very high speed and efficiency. In addition, it considers resilient and secure networking among all processing units in AI DCs. High-speed, low-latency communication among AI accelerators is primarily achieved through Remote Direct Memory Access (RDMA) fabrics allowing direct memory-to-memory data transfers between nodes without CPU involvement. Dominant implementations include RoCEv2/InfiniBand across racks in training clusters [14] and direct GPU-to-GPU communication within racks via NVLink [15].

Table 1
Hierarchical ESS classification framework for AI data centers.

Layer	Technology	Response timescale	Typical rating	Primary role	Section
Chip-level	On-chip capacitors, HBM buffers	Sub-ms to ms	W–kW	GPU power spike smoothing	7
Rack/Server-level	BBUs, supercapacitors, rack BESS	ms to seconds	kW–tens of kW	Server-level load buffering, BBU backup	7
Facility-level	Traditional UPS (TUPS)	Cycles (<20 ms)	100s kW–MW	Load protection, ride-through	3
	Grid-interactive UPS (GiUPS)	0.5–10 s (FFR)	100s kW–MW	FFR, VRT, peak shaving, frequency reg.	3
Grid-scale	BTM BESS	Minutes–hours	MW–100s MW	Peak shaving, RES integration, backup	4
	FTM BESS	Minutes–hours	MW–GW	Grid ancillary services, reserve	4
Grid-scale	Fuel Cell (FC)	Seconds–minutes	100s kW–MW	Low-carbon backup, co-generation	5
Facility-level	Thermal Energy Storage (TES)	Minutes–hours	MW thermal	Cooling load shifting, RES support	6
Facility-level	SLBESS	Seconds–hours	MW–100s MW	BTM/FTM backup, cost-effective BESS	8

2.1.1. *Computing units*

The computing unit consumes half of the required electricity of an AI DC [16]. NVIDIA GPUs such as H100, B200, or Blackwell Ultra have been used for training jobs in hyperscale AI DCs. For instance, the B200 utilizes 208 billion transistors that can be parallelized through very fast communication settings such as NVLink [17]. In addition, new GPU designs such as the GB300 NVL72 enable chip-level power smoothing using ESS [18]. Beyond NVIDIA, major cloud providers have developed custom AI accelerators optimized for their infrastructure: Google's Tensor Processing Units (TPUs) [19] are used to train and serve models such as Gemini, while AWS Trainium and Inferentia chips provide training and inference acceleration, respectively, enabling cloud-scale AI services.

AMD has also entered the AI accelerator market with its Instinct GPU series (e.g., MI300X), which offers competitive performance-per-watt for LLM training and inference workloads. These GPUs integrate CPU and GPU compute with high-bandwidth memory (HBM) in a unified package, positioning it as a direct competitor in hyperscale AI DCs. In Section 7, the role of ESSs for power smoothing will be discussed in more detail.

2.1.2. *Data storage units*

AI training workloads place extreme demands on storage infrastructure, requiring high-throughput, low-latency access to datasets that can reach petabyte scale. NVMe SSDs have become the dominant storage medium in modern AI DCs, replacing both HDDs and legacy SATA SSDs due to their superior bandwidth and latency characteristics. HDDs, while still present in archival or cold-storage tiers of some traditional facilities, are largely absent from the hot data path in AI training environments, where even brief I/O stalls can idle expensive GPU clusters [20].

The prevailing architectural trend in hyperscale AI DCs is *disaggregated storage*, in which compute (GPU servers) and storage are decoupled into independent resource pools interconnected via high-speed RDMA fabrics. Unlike traditional server-attached storage, disaggregation allows compute and storage capacity to scale independently, improves overall resource utilization, and enables storage to be shared across multiple training jobs simultaneously [14]. In practice, GPU servers access remote NVMe devices over RDMA-capable networks using protocols such as NVMe over Fabrics (NVMe-oF), achieving latencies that approach those of locally-attached storage. NVIDIA's GPUDirect Storage extends this further by enabling direct DMA transfers between NVMe storage devices and GPU memory, bypassing the CPU and host memory entirely and reducing both latency and CPU overhead for I/O-bound training workloads [21].

2.1.3. *Networking*

AI DCs aggregate thousands of GPU accelerators into tightly coupled parallel training clusters, placing extraordinary demands on the underlying network fabric. Unlike TDCs, where network traffic patterns are largely client–server in nature, AI training workloads generate dense all-to-all communication patterns during gradient synchronization across distributed model replicas. This requires network fabrics that deliver high bandwidth, ultra-low latency, and minimal congestion at scale [14].

The foundational transport mechanism in modern AI DC networks is RDMA, which enables direct memory-to-memory transfers between nodes without CPU involvement, reducing both latency and CPU overhead. While InfiniBand has historically been used in HPC settings and remains present in some training clusters, the industry has broadly converged on RoCEv2 (RDMA over Converged Ethernet) [22] as the fabric of choice at hyperscale, offering comparable performance over commodity Ethernet switching infrastructure at lower cost and with greater operational flexibility [14,23]. NVIDIA's Spectrum-X platform is specifically engineered to optimize RoCEv2 performance for AI workloads, addressing congestion control and load balancing challenges that arise in large-scale training deployments [23].

Within a single node or rack, inter-GPU communication bypasses the network fabric entirely through NVIDIA NVLink, a high-bandwidth direct interconnect that enables GPU-to-GPU data transfers at bandwidths far exceeding those achievable over any network interface [15]. Beyond a single rack, scale-out communication relies on the RDMA fabric described above.

2.2. *Power delivery systems*

The power delivery systems in an AI DC contain the external utility grid, internal power distribution, UPS, on-site backup generation, and the associated control systems.

2.2.1. *External utility grid*

TDCs, with power consumption on the order of a few megawatts, can typically be directly connected to distribution grids at voltage levels such as 34.5 kV or 13.2 kV. However, hyperscale AI DCs, requiring more than 100 MW of power, should be connected to transmission grids to ensure reliable and stable grid operation [2]. However, there is no standardized voltage level across utilities for integrating large loads such as AI DCs. The utility power is then transformed to a lower voltage level through a medium-voltage to low-voltage distribution transformer, typically 480 V, to feed the internal power distribution within the DC. This transformer also electrically isolates the AI DC from the grid. The 480 V AC is then distributed among the various components of the AI DC, including the UPS system, power distribution units (PDUs), cooling system, and IT equipment.

2.2.2. *UPS system*

The UPS system is comprised of an AC-to-DC power converter, a battery storage unit, and a DC-to-AC power converter. Power is stored in the storage unit under normal conditions, and the stored energy is used during emergency conditions when there is a loss of utility power or rapid fluctuations that jeopardize the normal operation of IT loads. In TDCs, UPS operation is typically limited to a few seconds, after which on-site generation units take responsibility for powering the DC in the absence of utility power for several hours.

The UPS requirements for AI DCs differ significantly from those of TDCs. The UPS system for AI DCs should operate over different ranges, including short-, medium-, and long-duration operation. Medium- and long-duration operation provides more time for on-site generation activation or, in some cases, can eliminate the need for separate backup units. GiUPS systems with high power density and large storage capacity can also be implemented to enable grid-support functions, including peak shaving or frequency regulation scenarios. The UPS systems in modern AI DCs will be discussed in detail later in Section 3.

2.2.3. Backup on-site generation

Backup generation units are utilized in DCs to provide electrical power under emergency conditions when utility power is unavailable. Backup generation is activated a few seconds after a utility power interruption, following the UPS response. Diesel backup generators are used in TDCs because they are combustion-based units that can provide a fast response; however, they are not environmentally friendly. RESs with appropriate BESSs can also be utilized for backup generation, as discussed in more detail in Section 4. However, power density, implementation cost, and the required space for installing PV panels or wind turbines are challenging issues in this context [24]. Modular nuclear reactors have also recently been introduced for AI DCs, offering MW-level power generation [25]. The technology for this type of equipment is still immature, implementation costs are high, and spent nuclear fuel must be stored properly. In addition, FCs may be better options for backup generation, as they can provide high power with fast response. Moreover, the development of hydrogen-based FCs can enable clean generation units within AI DCs [26]. The role of FCs in backup generation for future AI DCs will be discussed thoroughly in Section 5.

2.3. Cooling systems

The IT equipment inside a DC produces heat as a result of electrical losses within electronic chips. As the computational power and power density of computing units increase, the associated losses and heat generation also increase. In AI DCs, due to the presence of high-intensity AI accelerators, advanced cooling systems are required to enhance IT load performance and optimize power consumption. In TDCs, air is used as the primary medium for cooling IT equipment. Chilled air is injected into server rooms, and the temperature is maintained at a certain level using a central cooling controller. Air-cooling systems are cost-efficient, reliable, and agile. They can be easily extended to larger facilities, and the required hardware can be modified quickly. However, the efficiency of air-cooling systems is low for effectively cooling power-hungry AI accelerators.

Liquid cooling is a more advanced cooling system developed for AI DCs. By utilizing liquid cooling systems, cooling can be provided at the rack level or even at the chip level. Liquid cooling offers better heat transfer, requires less energy and space, and is more efficient compared to air-cooling systems. It also eliminates the need for large mechanical fans, which are less reliable. Direct-to-chip cooling circulates a dielectric liquid through cold plates into chips and processing units and can efficiently dissipate heat [27]. Hybrid cooling systems combine the agility of air cooling with the efficiency of liquid cooling, making them a suitable option for hyperscale AI DCs where redundancy can be achieved through hybrid solutions.

TE units can also be utilized in parallel with different types of cooling systems [28]. The TEs absorb chilled air or water and inject it into the IT equipment at specific times. Utilizing TEs can effectively increase efficiency in AI DCs. TEs cool the air or liquid by accessing extra energy from on-site generation units during normal DC operation. The role of TEs as cooling system accelerators in future AI DCs will be discussed thoroughly in Section 6.

2.4. Miscellaneous units

Apart from IT equipment, power delivery, and cooling systems, there are other units required for the normal operation of AI DCs. These include the physical infrastructure for housing servers, racks, and wiring systems, as well as the lighting system within the AI DCs. In addition, security and control rooms are also integral units of a AI DC. Miscellaneous units in an AI DC consume significantly less electricity than the other sections.

3. GiUPS systems in AI DCs

Building on the facility-level storage role introduced in Section 2, this section focuses on how UPS architectures evolve from passive backup devices into grid-interactive resources in AI DCs. TUPS architectures were designed primarily for load protection with minimal interaction with the utility grid [29–31]. This paradigm is increasingly insufficient for AI DCs. First, AI workloads introduce rapid load variations that propagate upstream and cause voltage flicker, transformer thermal stress, and feeder congestion [32]. These effects are amplified when multiple AI racks operate in synchrony, which is common in distributed training. Second, the growing penetration of renewable energy reduces system inertia and increases the rate of change of frequency. TUPS systems do not contribute to frequency support beyond simple ride-through, which leaves a large amount of inverter capacity unused during normal operation [33]. Third, the energy stored in UPS batteries remains idle for more than 99 percent of their lifetime, representing a significant opportunity cost. These limitations motivate the development of UPS architectures capable of dynamic grid interaction.

GiUPS systems enhance TUPS designs by incorporating bidirectional power converters, advanced digital controllers, and communication interfaces that allow coordination with the utility grid or market operator. These systems can modulate active power in response to grid frequency deviations or dispatch signals, enabling participation in services such as primary frequency regulation, fast frequency response, synthetic inertia, and peak load management [34–36]. Unlike TUPS systems, which operate independently of grid conditions except during outages, GiUPS systems continuously monitor grid conditions and adjust their output accordingly.

GiUPS can provide fast frequency response (FFR) after a grid-frequency event by detecting a frequency deviation or receiving a triggering signal and then increasing the battery power injection. In this paper, the detection or triggering time is distinguished from the full power output time: the former refers to the time associated with frequency measurement, event identification, and controller triggering, whereas the latter refers to the time required for the UPS battery to reach the specified FFR power setpoint. According to NERC, FFR is defined as active power injected into or absorbed from the grid in response to measured or observed frequency changes during the arresting phase of a frequency excursion event [37]. In the Nordic FFR technical requirements, the maximum time for full activation is specified as 0.7 s for the activation level of 49.5 Hz, 1 s for 49.6 Hz, and 1.3 s for 49.7 Hz, while the minimum support duration is specified as 5 s for short support duration and 30 s for long support duration [38].

GiUPS can operate in a dynamic frequency response mode, where the injected battery power is continuously adjusted according to the measured frequency deviation, frequency nadir, or rate-of-change-of-frequency requirement. It can also operate in a static frequency response mode, where a predetermined amount of power is injected once a frequency threshold or grid-event trigger is detected, such as a fault, generator outage, or large load switching event [39].

Table 2
Overview of GiUPS operating modes.

Mode	Grid input	Battery output	Trigger condition	Key feature
Standard Operation	100%	0%	Normal operation	Standard double-conversion UPS
Discharge-Full Disconnection	0%	100%	External command	Fully battery-powered load
Discharge-Partial Disconnection	<100%	Partial (e.g., 25%)	External command	Mixed grid and battery supply
Recharge	Supply > Load	Battery charging	SOC < 100%, over-frequency	Power limited (20%–25% of UPS capacity)
Energy Export	Adjustable	Discharge to grid	Regulatory approval	Bi-directional power conversion

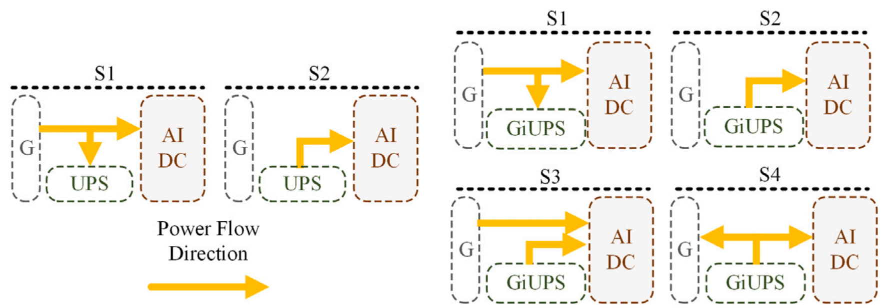

Fig. 2. UPS and GiUPS operating modes in AI data centers. The GiUPS extends conventional backup operation by enabling partial/full disconnection, recharge, and bidirectional grid-support functions while maintaining critical-load continuity.

3.1. Operational modes of GiUPS

GiUPS can operate in several modes depending on the grid conditions and external control signals, including *standard operation*, *discharge modes* (full and partial disconnection), *recharge mode*, and *energy export (bi-directional) mode* [34,40]. When a frequency variation is detected by an external controller, the GiUPS adjusts its operation to provide both positive and negative regulation by charging or discharging its batteries within operational limits. Table 2, shows an overview of GiUPS operation modes.

Under normal conditions, the GiUPS receives 100% of the input power from the utility grid through the rectifier, delivering it to the load (see S1 in Fig. 2). In this mode, the GiUPS functions as a standard double-conversion system, with the batteries maintained in standby and not actively supplying power. This operation is referred to as the *standard UPS operation*.

Discharge modes are employed to meet external control requests, supplying energy to the load or the grid. In the *full disconnection mode*, the GiUPS is completely disconnected from the utility, and 100% of the load power is supplied by the batteries (see S2 in Fig. 2). In the *partial disconnection mode*, the input power from the grid is reduced according to external commands, with the remaining load power supplied by the batteries (see S3 in Fig. 2). For example, the batteries may provide 25% of the total load power while the grid supplies the remainder.

When the battery state of charge (SOC) is below 100% (e.g., 80%) and an over-frequency event is detected, the GiUPS draws power from the grid to recharge the batteries. This operation is referred to as *recharge mode*. The maximum recharge power is typically limited to 20%–25% of the GiUPS nominal capacity, constrained by battery recharge characteristics and the maximum input current.

The GiUPS can also operate as a bi-directional power converter to inject energy back into the grid (see S4 in Fig. 2). This *energy export mode* is subject to local regulatory and grid requirements and allows the GiUPS to discharge batteries upstream.

3.2. GiUPS reactive power support and voltage ride-through (VRT) capability

In addition to active power modulation, advanced GiUPS systems can provide reactive power injection and VRT capabilities, which are increasingly required by modern grid codes and interconnection standards [41,42]. Reactive power support is enabled by the bidirectional inverter stage, which can independently regulate active and reactive current components under a synchronous reference frame control strategy. This functionality expands on traditional grid-support schemes where battery storage provides dynamic reactive power support and line-rating coordination to preserve system flexibility. [43] GiUPS units typically provide reactive power capability in the range of 20%–40% of their rated kVA. For example, a 1 MW/1.1 MVA GiUPS can supply approximately ±250–300 kVAr of dynamic reactive support. GiUPS can inject or absorb reactive power without significantly affecting active power delivery to the critical load. This capability allows the DC to contribute to feeder voltage regulation, mitigate voltage fluctuations caused by rapid AI load variations (often inducing 5%–10% voltage swings within milliseconds), and improve local power factor at the point of common coupling (PCC) from typical values of 0.92–0.95 to above 0.99 [41,44].

VRT capability ensures that the GiUPS remains connected and operational during short-duration voltage sags or swells [45]. TUPS systems typically transfer to battery mode during severe disturbances, effectively isolating the load from the grid. In contrast, GiUPS systems are designed to comply with low-voltage ride-through (LVRT) and high-voltage ride-through (HVRT) requirements, similar to grid-scale inverter-based resources. Typical LVRT profiles require the inverter to remain connected down to 0.2 pu for 100–150 ms, and down to 0.5 pu for up to 500 ms. During a voltage sag, the GiUPS can increase reactive current injection proportionally to the voltage deviation, often delivering 2–3 pu reactive current within 2–5 ms, thereby supporting grid voltage recovery while maintaining uninterrupted supply to the critical load.

The coordination between active power support, reactive power regulation, and VRT functionality requires a hierarchical control framework [46–48]. In grid-following (GFL) mode, the GiUPS synchronizes to the grid using a phase-locked loop (PLL) and provides reactive current according to predefined droop characteristics (typically 3%–5% voltage droop) or grid operator commands. In grid-forming (GFM) mode, the inverter can regulate terminal voltage magnitude and frequency directly, enabling seamless transition to islanded operation if grid conditions deteriorate beyond acceptable thresholds. Protection logic ensures that grid support functions never compromise critical load reliability, maintaining load voltage within ±1%–2% even during external disturbances.

3.2.1. GiUPS topologies

For GiUPS applications, topologies that support bidirectional power flow and high-speed digital control are preferred. Double-conversion and delta conversion architectures with bidirectional converters can rapidly transition between load-following, grid-support, and islanded operation. Modular architectures are particularly attractive in AI DCs because they allow some modules to participate in grid services while others remain dedicated to critical load protection. This partitioning reduces operational risk and improves economic performance [34]. Emerging DC-coupled architectures further reduce conversion stages and improve round-trip efficiency. These architectures also enable direct coupling with on-site renewable generation or FCs, which enhances the ability of the AI DCs to operate as controllable grid resource.

3.3. ESS sizing in GiUPS

GiUPS ESS sizing must balance two objectives: maintaining backup reliability and delivering grid services such as FFR, primary regulation, and peak shaving [11,12]. While backup needs establish the baseline energy capacity, grid-service participation adds important power-rating and cycling requirements, since services like frequency regulation are more power-intensive than energy-intensive. Economic sizing is therefore strongly influenced by cycle life, cumulative energy throughput, temperature, and average state of charge, because frequent shallow cycling can accelerate degradation [49–56]. Moderate oversizing can reduce effective DoD and extend battery lifetime, but excessive oversizing increases capital cost, footprint, and thermal management burden without proportional benefit. Accordingly, optimal sizing should jointly consider backup energy requirements, committed grid-service power capacity, and degradation-aware lifecycle cost modeling. For AI DCs, this process must also account for additional reserved capacity for short-term load smoothing, since this function helps mitigate upstream ramp rates and transformer stress but should not compromise guaranteed backup autonomy.

4. BESS in AI DCs

While GiUPS provides fast facility-level protection and grid interaction, dedicated BESSs extend this capability to longer-duration power smoothing, renewable integration, and ancillary-service support. Following the ESS hierarchy introduced earlier, this section focuses specifically on grid-scale and facility-level BESS functions that operate on minutes-to-hours timescales. Within the four-layer taxonomy, BESS occupies the facility- and grid-scale layers, complementing chip- and rack-level storage that handles sub-second transients. More specifically, this section reviews the fundamental BESS control paradigms, with emphasis on grid-forming and grid-following modes of operation. The discussion then extends to key grid-support functions enabled by AI DC-connected BESS, including co-generation, power smoothing, and frequency regulation. Practical deployment examples are included to ground the discussion in real-world infrastructure. Finally, the cost and sizing implications of BESS deployment are examined.

4.1. BESS control modes

From an operational perspective, AI DCs can be configured for grid-tied or islanded operation, depending on the BESS inverter control mode and the facility-level coordination architecture. GFM control enables the BESS to establish the voltage and frequency reference required for islanded operation. In grid-tied mode, the BESS can charge or discharge to support peak shaving and increase renewable self-consumption, whereas in islanded mode it serves as the anchor resource to supply the AI DCs and potentially reduce reliance on conventional generation or support renewable co-generation [57,58].

BESS inverters achieve this through one of two fundamental control paradigms. In GFM operation, the BESS behaves as a controlled voltage source that independently establishes its own terminal voltage magnitude and frequency without relying on an external phase-locked loop (PLL), emulating the inertial swing dynamics of a synchronous generator through an active power–frequency droop law, $\omega = \omega_0 - k_P(P_{\text{mes}} - P_{\text{ref}})$, and a reactive power–voltage droop, $V = V_0 - k_Q(Q_{\text{mes}} - Q_{\text{ref}})$ [59,60]. This self-synchronizing capability makes GFM the preferred paradigm for islanded AI DC microgrids or weakly connected sites, where it has been shown to suppress voltage flicker by up to 98.93% and maintain ride-through above the 0.9 pu threshold during fault events, keeping large AI DC loads continuously connected rather than tripping to local supply [59]. In contrast, GFL BESS operate as controlled current sources that synchronize to an externally imposed grid voltage via a PLL [61]; this mode is well suited to strongly meshed, grid-tied AI DCs where the primary objective is energy arbitrage or demand-charge reduction, but its performance degrades under weak-grid conditions as the external voltage reference becomes insufficiently stiff [62]. Hybrid GFM–GFL schemes that combine the voltage-forming stability of GFM with the current-limiting protections of GFL are therefore increasingly recommended for large AI DCs that must satisfy ride-through requirements across a wide range of grid-strength conditions [63].

4.2. Grid support roles of BESS in AI DCs

4.2.1. On-site clean power integration

BESS play a critical role in enabling RESs integration at AI DCs. Given the energy-intensive nature of AI DCs, on-site renewable generation and off-site procurement mechanisms such as power purchase agreements (PPAs) are increasingly adopted to reduce operating costs and meet sustainability objectives. However, the inherent intermittency of renewable resources complicates reliable power supply. In wind-integrated networks, battery storage paired with network topology or dynamic thermal rating optimizations has been established as an essential mechanism for maximizing renewable penetration and stabilizing system flexibility. Because AI data centers introduce a new class of highly volatile load behavior, extending these storage-enabled reliability frameworks becomes crucial for sustaining stable grid interaction under changing system conditions [64,65]. BESS mitigate these challenges by storing excess energy during periods of high generation and supplying it during low availability, while also providing backup power, and improved renewable utilization [66]. AI DC-integrated BESS can be architected either as large, multi-megawatt GFM resources capable of replacing conventional backup generators, or as large-capacity batteries integrated within static UPS systems [67]. In this context, GiUPS systems, traditionally designed for short-duration ride-through, are increasingly complemented by dedicated BESS that are optimized for extended energy management and grid-interactive operation rather than transient backup alone [58]. At the converter-control level, BESS can be configured for GFM, GFL, or hybrid GFM and GFL operation to balance fast dynamic response, voltage/frequency support, and grid-service capability, although with added control and sensing complexity.

4.2.2. Power smoothing

BESS can inject or absorb power at the facility interface to attenuate rapid AI DC demand variations and present a smoother power profile to the grid. A BESS co-located with an AI DC can provide grid-level power smoothing while supplying sufficient energy capacity to support sustained mitigation actions. This architecture improves operational reliability by enabling load shaping, demand flexibility, and extended ride-through for AI workloads.

As shown in Fig. 3, the GFM BESS effectively smooths the power fluctuations of the AI DCs resulting from training and inference jobs. An alternative approach is a hybrid E-STATCOM and BESS configuration, which combines the fast, high-speed charge/discharge response of supercapacitors with the higher energy capacity of BESS to achieve

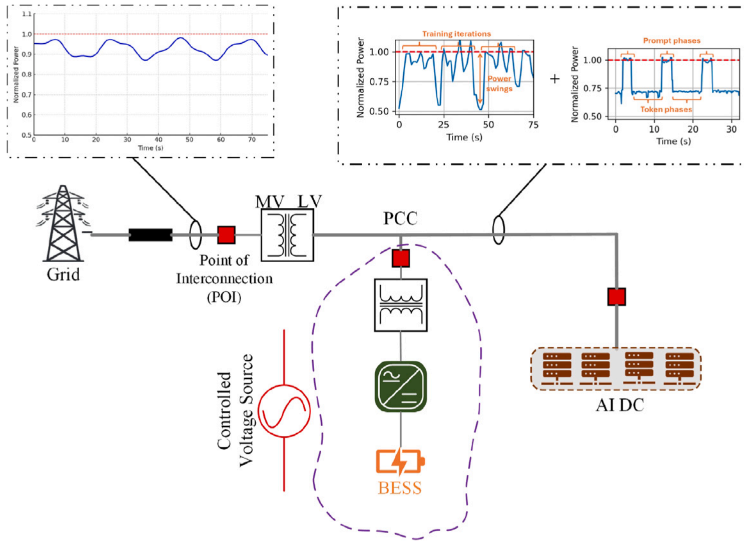

Fig. 3. Grid-forming BESS used as a power-smoothing unit for AI data centers. The BESS absorbs or injects power near the point of interconnection to reduce AI workload fluctuations seen by the utility grid [68].

effective smoothing and sustained load support. The real-time operational management of such configurations mirrors adaptive, multi-stage optimal dispatch frameworks that utilize deep reinforcement learning and real-time feedback to coordinate hybrid storage in active networks. Furthermore, from a grid-wide perspective, deploying high-capacity storage assets significantly limits power grid vulnerability and mitigates system stress [69,70]. This hybrid architecture can enhance resilience by improving demand flexibility, reducing flicker, and enabling broader grid-service capability. The primary trade-off is increased cost and footprint compared with standalone solutions. In addition, standalone supercapacitor-based E-STATCOM solutions can mitigate high-frequency AI load fluctuations and improve power quality for sensitive IT equipment by reducing flicker, although they provide limited long-duration energy support [68]. Table 3 compares the benefits and challenges of representative BESS-based and hybrid strategies for enabling grid support from AI DCs.

To illustrate the practical impact of BESS power smoothing, consider a 200 MW hyperscale AI training cluster transitioning from an idle checkpoint phase to full GPU utilization within approximately 200 ms—a ramp rate exceeding 600 MW/s at full scale. Without active smoothing, this ramp propagates directly to the transmission interface. A co-located 50 MW/100 MWh LFP BESS operating in GFM mode can absorb up to 50 MW of the transient, reducing the net grid-side ramp to below 150 MW/s and keeping the RoCoF at the PCC within the 0.5 Hz/s threshold specified by NERC reliability guidelines [37]. Field measurements reported in [68] confirm that properly sized GFM BESS can reduce peak-to-valley load excursions at the PCC by 40%–60% during AI training job transitions, substantially lowering reserve margin requirements for the serving utility.

4.2.3. *UPS-integrated BESS*

UPS-integrated BESS are motivated by a simple premise: AI DCs already deploy UPS systems with fast power-electronic interfaces, and augmenting this mandatory infrastructure with BESS enables grid-interactive capability with limited additional integration complexity. In practice, grid-interactive deployments also show that battery technology selection depends on the targeted services and operating constraints. For example, UPS-integrated BESS have been used to provide FFR by leveraging the rapid controllability of UPS power converters [67,72].

The UPS-integrated BESS, labeled as *power smoothing units* in Fig. 1, extend this concept by supporting long-duration outage operation and enabling participation in demand response programs, while the UPS maintains continuity of critical loads when neither the grid nor the BESS can fully supply demand. Although this architecture increases installation cost, it can significantly improve overall system reliability [35].

As summarized in Fig. 4, a GiUPS-integrated BESS can operate in (i) standard mode, (ii) discharge mode under full or partial disconnection, (iii) grid-services mode, and (iv) recharge mode for SoC recovery [34]. This architecture can reduce the need for bidirectional power flow across the MV/LV distribution transformer by localizing fast power exchanges, allowing the transformer to operate primarily under its normal unidirectional loading conditions [73]. These operating modes within the UPS–BESS subsystem enable long-duration BESS to enhance DC resilience through extended backup, reduce or eliminate reliance on diesel generation, participate in ancillary and reserve markets, manage demand charges and time-of-use tariffs, and increase RESs utilization [58].

(S1) Normal operation: The AI DC load is supplied entirely from the utility grid, and the system delivers 100% of the input power to the AI DC.

(S2) Full disconnection with sufficient BESS reserve: The site is fully isolated from the utility, and upon command from an external controller, the BESS supplies 100% of the AI DC load power. In this case, long-duration, large-scale BESS becomes essential; salt-cavern redox flow batteries are a promising due to their high safety, large storage capacity, stable temperature, and low cost [74].

During the transition from Scenario S1 to S2, the UPS initially supports the load because its response time is on the order of milliseconds, while the BESS typically requires a few seconds to assume full load supply. Once engaged, the BESS can sustain operation for extended durations (e.g., on the order of 1–4 h, depending on sizing and operating conditions) [58].

(S3) Full disconnection with insufficient BESS reserve: If the outage duration is long and the BESS does not have sufficient available capacity to supply the full AI DC load, the UPS assumes responsibility for supporting the critical load.

(S4) Partial disconnection: The external controller reduces the grid input power, and the remaining portion of the AI DC demand is supplied by the BESS. This operating mode enables the BESS to support grid frequency regulation and to present a smoother power profile at the grid interface.

(S5) Coordinated UPS–BESS power smoothing: The grid, BESS, and UPS jointly supply the AI DC load, with the UPS helping the BESS by providing fast support during critical transient conditions to assist

Table 3
Summary of power-smoothing architectures for AI data centers: benefits, challenges, and representative deployments.

Architecture	Response time	Benefits	Challenges
Advanced GFM	Seconds;	Load smoothing/shaping	Complex hybrid control requiring coordinated GFM/GFL mode transitions
+	FFR: 0.5–2 s;	Demand flexibility	Risk of active-power oscillations at the PCC if GFM and GFL units interact without proper decoupling filters
GFL BESS Solution	Peak shaving: minutes	Grid services Lower cost and footprint	Fast measurement and communication, <10 ms latency Advanced plant-level control
Representative example:	Microsoft–Eaton EnergyAware UPS pilot [67]: BESS co-located with UPS provides FFR to the grid while the GFM layer maintains load voltage.		
Hybrid E-STATCOM + BESS Solution	Sub-cycle; E-STATCOM: <1 ms; BESS buffer: minutes	Load smoothing and/or shaping Demand flexibility Grid services	Higher cost, ~1.3–1.5×standalone BESS Increased footprint Control interaction between supercapacitor and BESS converters if current-sharing references are not properly coordinated
		Lower flicker	May require inner-loop bandwidth separation exceeding one decade
Representative example:	Quanta Tech field demonstration [68]: 10 MVA E-STATCOM combined with a 20 MWh LFP BESS at a 100 MW AI DC campus reduced flicker index by 62% and peak grid-side ramp rate from 45 MW/s to under 8 MW/s.		
Grid-Forming BESS	Milliseconds; Inverter current loop;	Load smoothing/shaping Demand flexibility	May require a line reactor to limit fault current contribution Higher flicker during GPU-cluster synchronous start-ups unless ramp-rate limiting firmware is active
(Standalone)	Seconds for outer power loop	Support grid services/stability	Protection coordination complexity: GFM BESS may sustain fault current beyond inverter ratings Plant-level control required
Representative example:	Crusoe Energy/Redwood Materials [71]: Grid-forming SLBESS installed at an AI DC site provides black-start capability and smooths training-job power ramps.		
Supercapacitors	Sub-millisecond	High-speed response	Low energy density, 2–10 Wh/kg vs. 150–250 Wh/kg for Li-ion
(e.g., E-STATCOM)	to millisecond	Effective load smoothing/shaping	POI demand management only; cannot sustain longer outages
		Reduced flicker Near-unlimited cycle life, > 10^6 cycles, with negligible degradation	Does not offer GFM support Requires pairing with BESS or fuel cell for energy-duration requirements exceeding ~30 s
Representative example:	NVIDIA GB300 NVL72 rack-level supercapacitor buffer [18]: On-board capacitive storage at the GPU shelf level absorbs millisecond-scale power bursts preventing propagation to facility BESS and reducing required BESS power rating by up to 15%.		

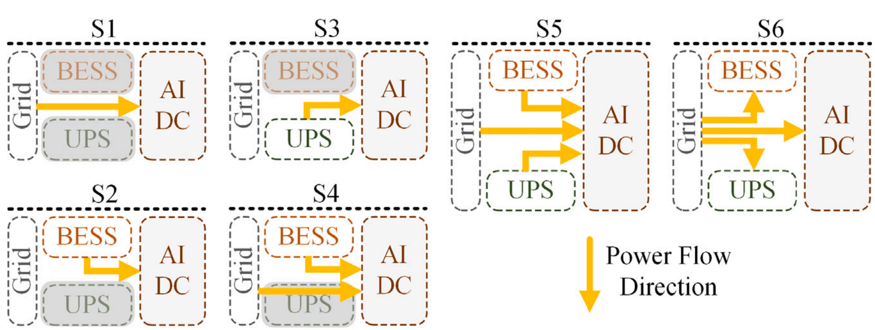

Fig. 4. UPS-integrated BESS operation modes for AI data centers. The coordinated UPS–BESS system supports normal operation, full/partial disconnection, transient smoothing, grid services, and recharge.

power smoothing and frequency response. This case can be interpreted as an extension of S4 in which the UPS contributes rapid, short-duration buffering when required.

(S6) Recharge mode: When the SOC drops below full charge, the utility initiates battery recharging (e.g., following an over-frequency detection). Under this condition, the UPS is prioritized for recharging first, after which the BESS can be recharged from the grid or from on-site renewable generation.

4.3. BESS cost and sizing considerations

The cost of BESS per kilowatt-hour has declined dramatically, dropping from approximately \$7500 in 1991 to \$133 in 2024, with projections indicating a further reduction to around \$80 by 2030 [58]. This downward cost trajectory is expected to continue with further scale-up in manufacturing capacity and advances in both manufacturing and battery recycling processes.

In addition to capital cost, effective sizing must account for operational objectives. Co-generation, power smoothing, and assisting GiUPS each require independent studies to determine the optimal BESS size. Sizing also affects the number of BESS units required for AI DCs. This aligns with broader power system methodologies where the optimum allocation and sizing of battery assets are leveraged to strengthen grid performance under variable operating conditions [75,76]. The sizing study will provide precise information on whether a single hyperscale BESS is more optimal than multiple grid-scale BESS units. Recent studies have applied advanced methods to determine the optimal BESS capacity while ensuring the required capacity and satisfying dispatch

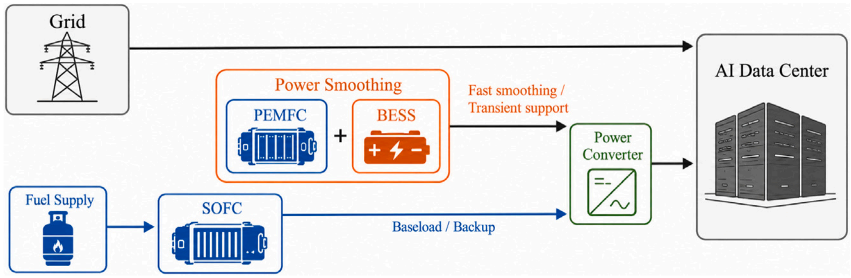

Fig. 5. FC integration architecture for AI DCs. SOFCs provide sustained baseload or backup generation, while PEMFCs can complement BESS-based smoothing under degradation-aware control.

constraints [77–79]. Ref. [80] shows that a 560 kWh Lithium-ion BESS, along with a 960 kWh lead–acid BESS, can provide FFR and backup power generation for TDCs. This study can be further extended to AI DCs as well, in which the cooperation between UPS and BESS is also considered.

5. Fuel cells (FCs) in AI DCs

While GiUPS and BESS architectures can provide fast ride-through, power smoothing, renewable integration, and grid-support services, they are still constrained by stored energy capacity, degradation, and cost when long-duration backup is required. For AI DCs with very large and reliability-critical loads, extended outages or sustained islanded operation may therefore require complementary on-site generation. In this context, FCs, through their diverse fuel conversion pathways, represent a promising low-carbon alternative to conventional diesel generators. Beyond sustained backup generation, FC-based systems may also contribute to slower facility-level power balancing when coordinated with fast-acting ESSs, power-electronic interfaces, and appropriate energy-management controls. However, FCs are not well suited to replace the technologies discussed previously for high-frequency GPU-level transient mitigation. In addition, many FC-based grid-support and control architectures remain at early stages of maturity, particularly for large-scale AI DC applications. Therefore, this section reviews the role of FCs in AI DCs from both a complementary generation and coordinated power-support perspective.

FCs convert chemical energy directly into electrical energy through electrochemical oxidation reactions. Because this conversion is not combustion-based, FCs can achieve electrical efficiencies exceeding 60%, which is significantly higher than many conventional generators [81]. FC systems offer high operational reliability due to their modular design and minimal use of mechanical components, resulting in reduced maintenance requirements and high availability. When supplied with hydrogen, FCs produce no carbon dioxide or conventional air pollutants during operation [82].

Among the several established FC chemistries, Solid Oxide Fuel Cells (SOFCs), in context of AI DCs, are often highlighted for their high steady-state efficiency and suitability for long-duration on-site power generation [83]. Polymer Electrolyte Membrane Fuel Cells (PEMFCs) exhibit faster electrochemical response and greater compatibility with power-electronic interfacing, making them promising candidates for coordinated power smoothing [84]. As illustrated in Fig. 5, both SOFC and PEMFC systems rely on a shared upstream fuel supply while serving distinct operational roles: SOFCs assist with base-load power to the AI DC, whereas PEMFCs, co-located with BESS, operate at faster timescales to mitigate short-term power fluctuations. With a consistent and reliable fuel supply, this architecture enables hierarchical coordination between the fuel infrastructure, on-site generation, and the utility grid.

5.1. Role of FC in AI DCs

Commercial deployments, such as Bloom Energy's SOFC systems, despite their relatively high capital cost, have demonstrated the feasibility of FC installations capable of supplying primary or supplemental power to TDCs. These systems are typically configured to provide continuous or long-duration power, serving as low-emission and energy-efficient alternatives to diesel generators [85]. In emerging AI DCs, whose power profiles are characterized by large, rapid, and highly synchronized load variations, FC-based technologies must be integrated within a broader energy architecture and operated as slow-following or baseload resources in order to preserve efficiency and lifetime. Accordingly, two complementary roles are considered:

5.1.1. FC on-site generation

SOFC-based systems are primarily suited for on-site generation in AI DC environments due to their high steady-state electrical efficiency, fuel flexibility, and scalability to megawatt-class power levels [83]. When supplied by pipeline natural gas or on-site hydrogen storage, SOFC systems can provide sustained power delivery over extended durations, making them alternatives to diesel generators for backup and resilience applications. Unlike battery-based systems, whose backup duration is limited by stored energy capacity, FC-based backup duration is primarily constrained by fuel availability, enabling outage coverage ranging from hours to multiple days depending on infrastructure configuration. However, due to their high operating temperatures and slow thermal dynamics, SOFC systems are not designed for rapid power modulation and are therefore best operated as baseload or slow-following generators within a coordinated energy management framework.

5.1.2. FC-assisted power smoothing

PEMFCs operate at significantly lower temperatures than SOFCs and exhibit faster electrochemical response, enabling limited power modulation relative to high-temperature FC technologies. This dynamic capability has motivated the use of PEMFCs for grid-support and ancillary service applications. Experimental work on a 5 kW PEMFC system showed that, under coordinated reactant-supply control, stack power could be increased from 2.0 kW to 4.0 kW within 0.1 s, corresponding to a 40% rated-power increase, or approximately 4 p.u./s on a 5 kW base [86]. These results demonstrate that PEMFCs can provide rapid power modulation at the stack or system level, but that their ramping capability depends strongly on reactant supply, auxiliary components, thermal management, and operating conditions.

Recent PEMFC grid-forming studies further show that ramp-rate limits must be explicitly considered to avoid fuel starvation. In one 120 kW PEMFC-based grid-forming system, the stack response time was modeled as 0.2 s, and the allowable ramp rate at the maximum-efficiency operating point of 55% was limited to 50 kW/s. When operated at lower efficiency, the allowable ramp rate increased to 120 kW/s at 50% efficiency and 500 kW/s at 45% efficiency [84]. These values illustrate an efficiency-ramping tradeoff: operating closer to maximum efficiency improves steady-state performance but reduces the allowable rate of power change, whereas faster ramping may require lower-efficiency operation, excess reactant supply, or additional energy buffering.

For AI DCs, these reported ramping results should be interpreted as evidence of PEMFC dynamic capability rather than as direct proof of standalone AI-load smoothing at hyperscale. PEMFCs are better suited to following ramp-rate-limited, low-frequency facility-level power references, while GiUPS systems, BESSs, supercapacitors, and rack-level buffers regulate millisecond-to-second transients and maintain AI DC-bus stability. When coordinated in this manner, PEMFCs can complement BESS-dominated smoothing strategies by supplying average power components and extending system endurance without being directly exposed to damaging high-frequency AI workload variations [87].

5.2. Sizing considerations

Scaling FC technologies to AI DC capacities introduces constraints related to installed power rating, hydrogen logistics, balance-of-plant complexity, and the temporal structure of the load assigned. Thus, FC sizing should be linked not only to required power and backup duration, but also to allowable ramp rate

For SOFC-based generation, the relevant sizing basis is the steady or slowly varying portion of AI DC demand. SOFCs are well matched to baseload supply, slow-following operation, and long-duration backup, but their high operating temperature and thermal inertia limit rapid power modulation. Their backup duration is primarily determined by fuel availability rather than stored electrochemical energy, making hydrogen storage, pipeline supply, reforming infrastructure, and redundancy central. Hybrid SOFC systems with thermal bottoming cycles have demonstrated stable operation at multi-megawatt and hundred-megawatt scales [83]; however, increasing capacity shifts the main design burden toward thermal management, mechanical integration, and fuel-processing infrastructure.

For PEMFC-based power support, the critical sizing constraint is the mismatch between AI DC load-ramp rates and allowable FC ramp rates. Since existing PEMFC grid-support demonstrations remain far smaller than hyperscale AI DC load variations, direct extrapolation is not sufficient. Multi-megawatt deployment would require parallel stack aggregation, coordinated hydrogen distribution, auxiliary subsystem synchronization, and high-capacity DC/DC and DC/AC conversion stages [88]. The PEMFC reference should therefore be filtered so that the requested trajectory remains within fuel-supply, air-supply, thermal, and degradation limits. In a hierarchical design, the PEMFC would follow the low-frequency component of facility demand, while faster ESS layers would handle the high-frequency component that exceeds the allowable ramp rate of the cell.

Although FCs can support long-duration electrical supply and low-carbon backup generation, they do not directly address cooling as another major source of AI DC energy demand. Because the heat generated by dense GPU clusters creates substantial cooling-related electrical load, flexibility on the thermal side can also reduce grid stress and improve overall energy efficiency. Therefore, after discussing electrical backup and on-site generation through FCs, the next section examines thermal energy storage as a complementary solution for shifting cooling demand and coordinating AI DC operation with grid and renewable-energy conditions.

6. Thermal energy storage (TES) in AI DCs

TES refers to technologies that store energy in the form of heat or cold by raising or lowering the temperature of a storage medium relative to ambient conditions [89]. Unlike electrochemical energy storage, TES does not directly support the electrical power and therefore cannot mitigate GPU-level electrical transients. Instead it reshapes the energy demand by decoupling thermal energy production from its consumption. This makes TES relevant in systems where thermal load are tightly coupled with performance and electrical demand, such as DCs [90].

In AI DCs, cooling constitutes 30%–40% of total facility energy consumption [91] due to the high power density of accelerator-based computing. Nearly all electrical energy consumed by GPUs and associated IT equipment is ultimately converted into low-grade heat, with studies estimating that around 90% of input electrical power manifests as waste heat within the DC environment [91]. Cooling demand, as a result, scales with computational load. While training DC power exhibits high-frequency power swings, thermal load evolves on slower timescales. This creates favorable conditions for the use of thermal energy storage to store cooling capacity and supply it during periods of high electrical demand or grid stress.

6.1. Characteristics and timescale of TES operation

TES exhibits operational characteristics that differ fundamentally from electrical ESSs. Due to thermal inertia and heat transfer dynamics, these systems respond on timescales typically ranging from several minutes to hours. As a result, TES is not suited for fast grid services such as frequency regulation or power quality support [92].

Instead, the storage units are well matched to applications involving load shifting, peak demand reduction, and ramp-rate mitigation of cooling-related electrical loads. Thermal standby losses and insulation constraints further limit the suitability of TES for long-term or seasonal storage in most data center applications. Consequently, TES is most effective when operated on diurnal timescales, complementing faster electrical storage technologies within a hierarchical energy storage framework [90].

6.2. TES technologies and integration architectures

6.2.1. Sensible TES

Sensible TES stores cooling capacity through temperature changes in a storage medium, most commonly water. Chilled water storage tanks are a mature and widely deployed form of the technology due to their simplicity, low cost, and ease of integration with conventional chilled-water systems [93]. However, because energy is stored through temperature differential rather than phase change, sensible TES exhibits comparatively lower volumetric energy density than alternative technologies. As a result, achieving equivalent storage capacity requires significantly larger physical volume and installation area.

Despite this footprint requirement, sensible water-based systems remain the least expensive option per unit of stored thermal energy and demonstrate predictable thermal behavior with minimal operational complexity [92]. For AI DC campuses where land availability permits, sensible TES provides a cost-effective and technically robust solution for short-term load shifting and peak chiller reduction.

6.2.2. Latent TES (PCM-based)

Latent TES exploits the phase change of materials to store thermal energy at nearly constant temperature. Compared to sensible storage, phase change materials (PCMs) provide higher volumetric energy density, allowing equivalent cooling capacity to be achieved with substantially reduced installation volume. This characteristic is particularly advantageous in space-constrained AI DC facilities [93].

The most widespread latent storage medium in cooling networks is ice, owing to its high latent heat and relatively low material cost. Ice-based TES has been successfully deployed in large-scale district cooling systems worldwide, demonstrating technical feasibility at substantial capacity [93]. However, latent systems introduce additional engineering complexity, particularly in heat exchanger integration, phase stability, and long-term cycling reliability. Although more compact than sensible storage, latent TES generally entails higher capital cost per unit of stored energy. Consequently, while ice-based TES offers clear volumetric advantages for cooling applications, its deployment in AI DC environments must balance footprint reduction against increased system complexity and economic uncertainty.

Table 4
High-level comparison of rack-level BBUs and centralized UPS systems.

Aspect	Rack BBUs	Central UPS
Primary purpose	Rack ride-through; local fault containment [98]	Facility ride-through; centralized management [40]
Timescales	Seconds–minutes (backup, recovery) [100]	Minutes+ (backup) and grid services where allowed [101]
Sub-second smoothing	Not ideal; accelerates battery aging [87,102]	Not ideal; GPU transients are localized rail events [87]
Coordination burden	High; fleet SoC, recharge scheduling, priority handling [103]	Lower; centralized components
Efficiency	Avoids some double conversion [104]	Often involves AC–DC–AC in classic designs

6.2.3. Thermochemical TES

Thermochemical TES stores energy through reversible chemical or sorption processes, allowing thermal energy to be stored as chemical potential rather than as temperature change. Such systems can achieve high energy density and the lowest standby losses, making them attractive for long-duration storage [93].

Despite these advantages, thermochemical TES systems are generally complex, expensive, and less mature than sensible or latent storage technologies. Their application in DCs is currently limited to research and demonstration projects, though they represent a potential future option as cooling demands continue to increase.

6.3. TES for RESs integration

TES can support the integration of RESs in AI DCs by enabling power-to-cooling strategies. During periods of surplus RESs generation or low electricity prices, chillers can be operated to charge TES systems, storing cooling capacity for later use [94]. This approach improves the utilization of RESs and reduces reliance on grid power during peak or high-emission periods.

By shifting cooling-related electrical demand in time, TES enhances the flexibility of AI DCs as large grid-connected loads and supports broader decarbonization objectives.

The previous sections focused on facility-level and complementary energy resources that support AI DCs through grid interaction, long-duration generation, and cooling-load flexibility. However, many AI workload fluctuations originate much closer to the IT load, particularly at the GPU, server, and rack levels. These fast local transients cannot always be managed efficiently by GiUPS, BESS, FC, or TES systems alone. Therefore, the next section shifts the discussion from facility-scale resources to rack- and server-level ESSs, which provide localized buffering near the source of AI power variability.

7. Rack- and server-level ESSs roles in AI DCs

Rack- and server-level ESSs exist to handle two constraints that centralized UPS/BESS alone cannot satisfy: (i) *local buffering* (GPU-/server-scale transients that require local energy buffering), and (ii) *fault domain* (limiting the blast radius of power disturbances and energy storage failures). In modern DC-architecture AI facilities, upstream layers (SST/MVDC front-ends and facility UPS/BESS) primarily set facility envelopes and ride-through at longer timescales, while downstream layers move buffering closer to the load to smoothen fast dynamics and localize contingencies under faults [95,96]. This section focuses on rack-level BBUs as distributed ride-through and backup resources, and on server/GPU-level capacitive storage as power profile smoothening for the highest-frequency behaviors [97–99].

In the four-layer hierarchy, rack- and server-level ESSs should not be viewed as storage devices operating in isolation. Their effectiveness depends on local control interfaces that enable the physical storage layer to address AI workload variability. At the chip and server levels, firmware and platform management interfaces can constrain GPU power ramps and expose power telemetry. At the rack level, BBU controllers manage SoC, discharge availability, and recharge scheduling. These control interfaces allow fast GPU-side fluctuations, rack-level smoothening and ride-through, and facility and grid-level UPS and BESS operation to be coordinated across timescales.

7.1. Rack-level ESS: Role, benefits, and constraints

The rack represents a natural operational and electrical boundary for AI infrastructure, with dense DC-architecture racks alone reaching MW-scale power demands [95]. Deploying energy storage at the rack level offers several advantages. First, the storage is electrically close to the load, enabling local buffering before propagating upstream to the facility distribution system. Second, failures in battery, converters, or wiring can be contained to individual racks, limiting the blast radius of faults. Third, storage capacity budgets and ride-through priorities can be configured on a per-rack basis, allowing workload-specific policies that reflect criticality differences across the DC.

Despite these benefits, rack-level deployment introduces operational challenges. Coordination becomes necessary, as many independent storage units must be managed for state-of-charge tracking, recharge scheduling, and availability monitoring. Safety and compliance requirements increase, since colocation of lithium-ion energy storage with IT equipment raises fire and thermal-runaway mitigation concerns. Finally, maintenance overhead grows substantially, as consistent monitoring, battery replacement, and preventive maintenance must be performed across racks. Fig. 6 shows the implementation of rack-level BBUs and server-level capacitors in the IT equipment.

7.2. Battery backup units (BBUs) as distributed UPS

In OCP Open Rack designs, BBUs are integrated into the rack as distributed DC ride-through modules that supply the rack DC bus during upstream AC disturbances [97]. In Open Rack V3 (ORV3) architecture, a BBU shelf typically hosts multiple BBU modules with 5+1 redundancy and supports both *charge mode* and *discharge mode* operations with monitoring of SoC/SoH and maintenance tests [98]. Public reference designs summarize the intended operating point (e.g., per-module backup power on the order of kW for minutes, and lower-power charging over hours), reflecting a design goal of short ride-through and fast recovery rather than continuous power profile smoothening [100]. Meta reports ORV3 BBU shelves designed for minutes of backup, and notes that paired shelves can be used for higher rack power configurations [105].

7.2.1. BBU vs centralized UPS

BBUs and centralized UPS systems can both respond quickly in principle to supply the gap between grid power disturbance or outage and longer-term energy storage systems coming online (e.g., gensets/BESS), but they differ in *where* energy is buffered and *what* the buffer is optimized for. Centralized UPS concentrates energy and power conversion at room/facility scale, simplifying management and enabling grid-interactive use cases when permitted [40,101]. Rack BBUs push storage behind fewer conversion stages in DC racks and reduce some double-conversion penalties compared to traditional centralized UPS architectures [104]. Practically, BBUs are best viewed as a distributed energy availability layer for the rack (seconds–minutes), not as the primary solution for sub-second smoothing which is handled more effectively by server/GPU capacitive buffers. Table 4 summarizes the high level comparison of rack-level BBUs and centralized UPS systems.

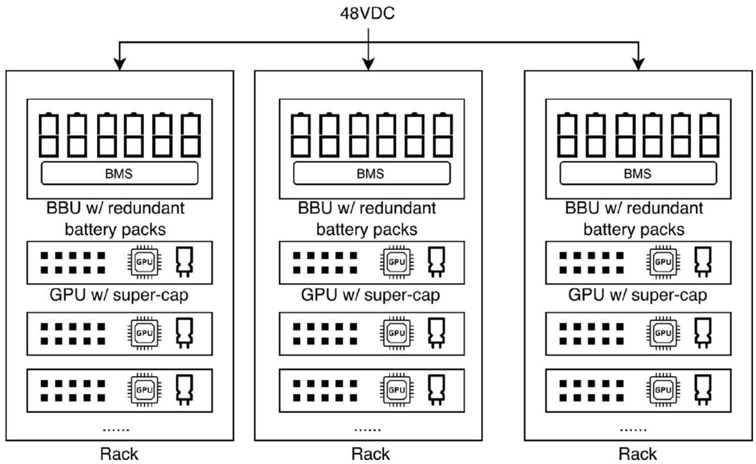

Fig. 6. Rack-level BBUs and server-level capacitors for localized energy buffering. BBUs provide rack ride-through, while capacitors suppress fast GPU/server power transients before they propagate upstream.

7.2.2. BBUs as power absorbers

BBU shelves inherently include charging control; in normal operation they remain in charge mode and can modulate charging current/power within design limits [97]. At hyperscale, uncontrolled simultaneous recharge can create step increases in aggregate load that trip upstream protection. Production case studies show that coordinated and priority-aware charging of distributed batteries can dramatically reduce recharge power (reported reductions up to ∼80%) while meeting recharge constraints [103].

However, using BBUs as a frequent, high-rate smoothing solution can accelerate degradation depending on cycling throughput, C-rate, SoC range, and temperature [102,106]. Practically, server/GPU capacitors and software solution are better suited to handle high-frequency power fluctuation, and BBUs are better suited for ride-through, controlled recovery, and low-frequency shaping with budgeted cycling.

7.2.3. Safety and deployment considerations

Rack BBUs introduce Li-ion safety considerations near IT equipment. OCP BBU specifications explicitly reference safety and propagation constraints and relevant certification/testing practices [97,107]. This strongly motivates designs that (i) enforce conservative SoC windows and thermal monitoring, and (ii) reduce unnecessary cycling.

If the design intent is “UPS-like” ride-through for all computing devices, deployment of BBU on every rack is the straightforward approach used in OCP-style architectures [97,98,105]. Selective deployment (e.g., BBUs only on the highest-power or highest-priority racks) can reduce capex and coordination overhead, but it changes the availability model: unprotected racks become dependent on workload-level fault tolerance, or upstream ESS (BESS/gensets) without transient protection.

7.3. Server- and GPU-level storage

7.3.1. Capacitors instead of batteries

The highest-frequency components of AI load power dynamics are at server-level: rapid power consumption swings on the GPUs. These events are handled by local decoupling and bulk capacitance [108]. Buffering at this layer prevents microsecond-to-millisecond disturbances from propagating to the power sharing unit (PSU) and rack bus, reducing upstream stress and allowing higher utilization at the rack power infrastructure [87].

Compared with batteries, supercapacitors tolerate extremely frequent charge–discharge cycles with minimal degradation, making them attractive for sub-second smoothing [109,110]. Industry offerings and proposals include compact supercapacitor banks intended to suppress short spikes that would otherwise require overbuilding upstream power infrastructure. The main constraint is energy density and physical volume: supercapacitors can buffer transients and short bursts, but they cannot provide minutes of ride-through and therefore complement (not replace) rack BBUs and other upstream ESSs [110].

7.3.2. Firmware and platform power control

Electrical buffering at the chip and server levels is increasingly paired with firmware-level power shaping. Rather than adding stored energy, these controls modify the load profile seen by upstream ESS layers. Instead of adding energy storage devices, these controls modify the load profile seen by upstream ESS layers. For example, NVIDIA describes firmware-controlled ramp-up behavior and a “power burner” mode to manage ramp-down and stabilize facility-level power dynamics during AI job transitions [18]. At the platform level, software interfaces such as NVIDIA’s NVML expose enforced GPU power limits and real-time power telemetry, enabling operators and higher-level controllers to constrain GPU power draw and coordinate device-level consumption with rack- or facility-level power envelopes [111]. It is framed as an industry direction to approach power stabilization as a cross-stack problem that couples hardware energy storage and software control to prevent large excursions and meet grid/facility ramp constraints [108].

7.4. Sizing considerations for rack- and server-level ESSs

Sizing at the rack and server layers is driven by different objectives than grid-scale or UPS-level storage. Rather than energy capacity for extended outages or grid-service participation, the dominant constraints are rack power consumption, backup duration target, and redundancy requirements.

The fundamental BBU shelf sizing requirement is that the shelf must supply the full rack power demand for the target ride-through duration under a single module failure. For AI racks, where power demands can exceed 100 kW and reach MW-scale in dense configurations [95], this places a substantially larger energy requirement on the shelf than in traditional deployments. In ORV3 architectures, the 5+1 redundancy model requires each active module to be sized for its proportional share of the full rack load plus margin [98]. For target backup duration, OCP reference designs reflect a goal of minutes of ride-through [100,105]. The specific duration that rack-level storage supports should be designed in sync with upstream ESSs, namely BESS or gensets, since its task is to bridge outage onset with extended ESS being online. Extending this window increases energy capacity but consumes already limited rack space, as well as amplifying thermal risks for Li-ion cells colocated with IT equipment [97,102].

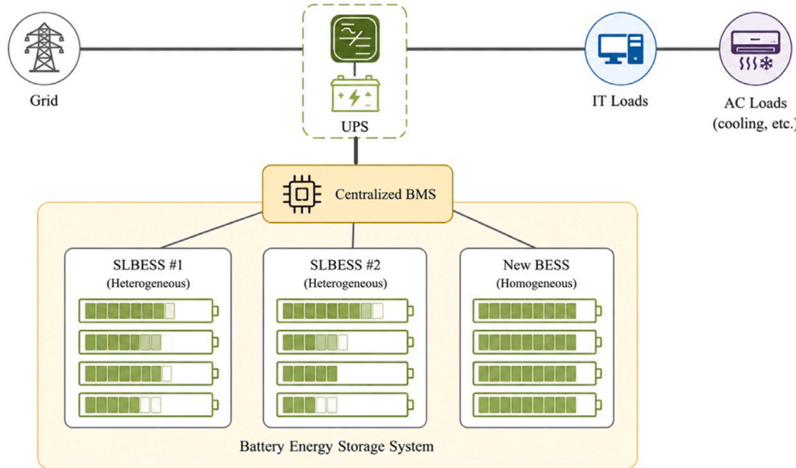

Fig. 7. SLBESS configuration for UPS-fed AI data center backup support. Heterogeneous second-life battery packs require centralized monitoring and health-aware control to ensure safe and reliable operation.

Supercapacitors at the server level are sized to absorb the transient energy of rapid GPU power swings [87]. For sub-second transients, while power swings could be large, the energy required is limited, which makes supercapacitors viable at this layer despite their low energy density [108,109]. The binding constraint is therefore capacity with regard to the GPU power swing rather than cycle life, since supercapacitors tolerate very high cycle counts with minimal degradation [18,110].

7.5. Coordination across rack and server layers

Coordination should be explicit about objectives and timescales: GPU/server controllers shape fast dynamics under performance constraints; rack BBU controllers manage SoC availability and recovery without creating recharge spikes at facility level; facility controllers enforce facility envelopes and power-quality limits [103,112]. OCP rack ecosystems already expose the necessary hooks (telemetry, shelf controllers, parallel shelf operation) to implement multi-layer coordination in practice [97,98]. Under MVDC architectures, server/GPU capacitive buffering can reduce required BBU power bandwidth and cycling, but does not eliminate the seconds–minutes energy role of BBUs for ride-through and controlled recovery [96].

Rack- and server-level ESSs complete the hierarchical view by showing how storage can be deployed close to AI computing loads to manage fast transients and localized reliability requirements. However, broader deployment of battery-based ESSs across AI DC layers also raises important cost, lifecycle, degradation, and sustainability concerns. These concerns are especially relevant for large facility-level and grid-scale storage systems, where battery replacement cost and material reuse can strongly affect practical adoption. Accordingly, the next section discusses SLBESSs as a cost-oriented pathway for expanding battery storage deployment in AI DCs.

8. Second-life battery energy storage systems (SLBESS) in AI DCs

Having established the functional roles of ESSs across architectural layers, an important practical question concerns cost scalability, motivating the use of second-life batteries (SLBs). Second-life battery energy storage(SLBESS) is a cost-effective application that can be deployed at multiple layers of the proposed hierarchy. Most current deployments are found at the facility and grid layers due to energy density, cost, and economic considerations.

While the application of SLBESS for grid support is well documented, their deployment for AI DC loads is an emerging and promising opportunity, as highlighted in a recent project by Crusoe Energy, where the company installed used batteries to power AI DC load [71]. SLBESS installations must also adhere to battery repurposing standards such as UL 1974 [113]. Therefore, regulatory compliance is mandatory for building AI DC infrastructure and battery reuse.

8.1. SLBESS assisting AI DC on-site generation

For AI DC load, SLBESS can serve as BTM generation backup source, provided that the battery packs are rigorously tested and qualified via health metrics. This is a GFL mode of operation where hybrid configurations, combining used and new battery packs, are technically feasible as shown in Fig. 7. The battery management system (BMS) balances the varying capacities of different packs. Additionally, a single SLB pack exhibits heterogeneity across cells and modules, resulting in variable capacities, voltages, and current limits. E.g. a single degraded cell may constrain the performance of an entire module or pack. To solve this, passive and active balancing schemes have been explored to mitigate these challenges. Advanced control architectures have been proposed to actively balance heterogeneous cell capacities in SLBs [114].

8.2. Hybrid GiUPS configuration for AI DCs using SLBESS

UPS systems in AI DCs provide instantaneous backup for critical IT loads. This necessitates energy storage solutions capable of sustained discharge durations exceeding 10 h. SLBESS are generally not suitable for extended durations, and operate within a narrow SoC window to avert deep discharge. This is because the over-discharge accelerates cell-level degradation and may push cells to the non-linear “knee point” region in the degradation curve. Thus, unless carefully derated and managed, SLBESS may offer limited suitability for long-duration UPS applications.

In AI DC, using SLBESS for short-duration backups while avoiding rapid degradation is technically feasible. Optimal dispatch strategies must be implemented to achieve that. Traditionally, the fundamental approach of using BESS involves energy arbitrage, where the BESS charging occurs when electricity prices are low, and discharge occurs when the electricity prices are high. This economic objective is balanced against the physical aging of the battery through a degradation cost component, shown in Eq. (1).

$$\max\left(C_{dis,t}\,P_{dis,t} - C_{ch,t}\,P_{ch,t}\right) - C_{deg}\left(\text{cycles}\left(P_{dis,t}, P_{ch,t}\right)\right) \tag{1}$$

subject to:

- Battery Constraints that include limits on the SoC and maximum charge/discharge rates to ensure operational safety.

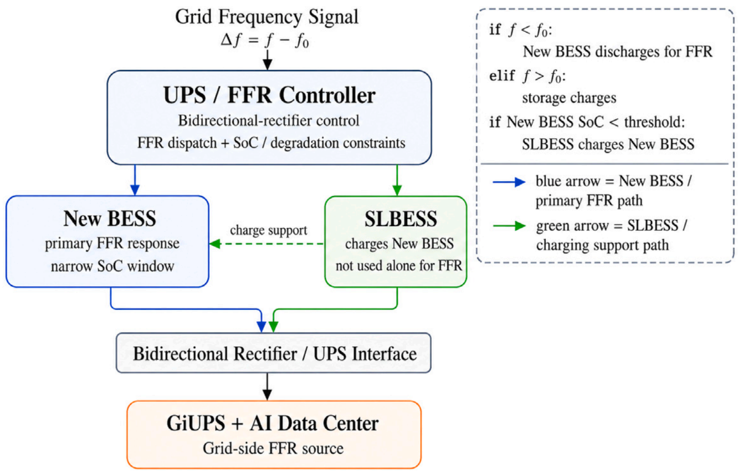

Fig. 8. FFR using coordinated new BESS and SLBESS resources in a GiUPS architecture. The storage units charge or discharge according to frequency deviations while respecting SoC and degradation constraints.

- Grid Constraints that includes power exchange limits and local grid regulations.

For ever time step, t in hours, C_{dis} and C_{ch} are the market price (\$/kWh) of electricity during the discharge and charge phase. P_{dis} and P_{ch} are the power output (kW) discharged from and charged to the SLBESS. C_{deg} is the degradation cost coefficient, representing the cost of battery wear per cycle. The cycles of BESS can be calculated as a function of aggregated energy throughput $f(P_{dis}, P_{ch})$. The sum of P_{ch} and P_{dis} represents the total power throughput, which serves as a proxy for the stress placed on the SLBESS. To find C_{deg} in \$/cycle, the data from an accelerated life testing of SLB cells of corresponding chemistry can be used.

For AI DCs, the primary objective of energy storage optimal dispatch shifts from pure arbitrage to supply guarantee and reliability. In a hybrid GiUPS configuration, the SLBESS must prioritize the load over economic gains. The optimality equation can be modified to include a reliability term, which serves as a significant penalty for unserved load, ensuring the AI DC remains operational. Therefore, Eq. (1) can be modified as Eq. (2)

$$\max\left(\begin{aligned}&C_{\mathrm{dis},t}P_{\mathrm{dis},t} - C_{\mathrm{ch},t}P_{\mathrm{ch},t}\\ &- C_{\mathrm{deg}}\left(P_{\mathrm{ch},t} + P_{\mathrm{dis},t}\right) - C_{\mathrm{reliability}}L_{\mathrm{unserved},t}\end{aligned}\right) \tag{2}$$

$C_{reliability}$ is a high-value penalty factor associated with the loss of power to the AI DC (interpreting the “supply guarantee” requirement from the sources) and $L_{unserved}$ is the portion of the critical AI DC load not met by the grid or the SLBESS while subjected to the BESS and UPS constraints.

On the grid side, GiUPS can be a source of FFR. This is possible with bidirectional rectifiers that can respond to changing grid frequency. FFR tends to use the BESS frequently within a narrow SoC window. SLBESS alone is not suitable for this sort of grid service, leading to its rapid degradation [115]. However, if used alongside a new BESS, the UPS controller can be programmed to leverage the SLBESS for charging the new BESS when its SoC falls below a certain threshold. If optimized, this can reduce the energy import from the grid, saving the cost of energy. Fig. 8 shows the possible scenarios, which are listed below.

(i) In case of grid frequency above the threshold, the UPS controller charges the new BESS and SLBESS, providing stabilization of the frequency.
(ii) In case of grid frequency below the threshold,
(a) The UPS controller can discharge new BESS and provide power back to the grid while powering the IT load with SLBESS.
(b)The UPS controller can meet a part of the IT load using the new BESS, and use SLBESS to charge the new BESS and stop any import/export to the grid.

8.3. Grid-scale SLBESS for AI DCs

The US national grid has a long queue for interconnection requests. This leads to limited generation and difficulties in accommodating new AI DC load. In [116], the authors concluded that 76 GW of new load can be integrated without significant upgrades by curtailing 0.25% of new flexible load per year. This amounts to a curtailment of 22 hours/year. FTM grid-scale SLBESS can achieve this by shaving peaks of AI DC load, decreasing the stress on the grid that is caused by erratic peaks of AI DC load. It can also acts as standby generation and partake in resource adequacy. Furthermore, as hyperscalers such as Meta, Microsoft, and Apple Inc. increasingly engage in energy market operations [117], grid-scale FTM SLBESS can be utilized to manage surplus power from renewable generation plants under PPAs enabling these hyperscalers to hedge against increasing energy prices.

By implementing the optimal dispatch strategies previously discussed, these systems can store excess renewable energy during periods of high production or low market cost and sell it back to the wholesale market when prices are high. This strategy is particularly effective when using lithium-ion batteries, such as lithium iron phosphate (LFP), which have been proven to be a viable and high-capacity option for DC loads.

However, to successfully integrate this FTM SLBESS approach into hyperscale infrastructure, several factors must be carefully balanced:

- **Reliability vs. Profit:** While selling energy to the wholesale market provides economic benefits, the dispatch logic must prioritize supply guarantees and the reliability of the AI DC load over arbitrage profits.
- **Degradation Management:** Because these systems utilize second-life battery cells, dispatch strategies must be carefully managed, and aging of the SLBESS must be monitored more rigorously than conventional BESS to prevent accelerated degradation during market participation.
- **System Optimization:** Hyperscalers can leverage real-time SLBESS status to optimize hybrid BESS/UPS system operation, enabling effective peak shaving and grid interaction while maintaining sufficient power reserves for critical AI DC operations.

8.4. Cost and sizing aspects of SLBESS for AI DCs

8.4.1. Csot

SLBESS is typically more cost-effective than new BESS. However, cost is influenced by supply, location, and market conditions. The cost of stationary storage declined from \$169/kWh in 2020 to about

$108/kWh in 2025, amounting to 36% reduction worldwide [118]. This decline trajectory may be a hindrance to SLBESS usage in the future. Moreover, refurbishing, operation, and maintenance can constitute a significant portion of CapEx. Market analysis suggests LFP batteries are more economically advantageous for second-life use compared to nickel-rich chemistries, primarily due to the NCX battery's higher recycling value posing a demand competition between SLB repurposers and recyclers [119]. Since reliability can be a huge factor in AI DCs, the hyperscalers may prioritize new BESS over SLBESS. However, the retirement of EV batteries is expected to coincide with rising DC energy demand and the policies, such as the *One, Big, Beautiful Bill Act* (OBBBA) bill, introduce *investment tax credits* (ITC) for battery storage sourced within the country, enhancing the attractiveness of domestically available SLBs for AI DC deployment [120].

8.4.2. Sizing

Sizing SLBESS for AI DCs presents a distinct engineering challenge. SLBs, repurposed from EV applications, exhibit reduced usable capacity, increased internal resistance, and greater heterogeneity in state-of-health (SoH) compared to new cells [121,122]. To mitigate accelerated degradation and ensure operational reliability, SLBs are typically operated within narrower SoC windows and under derated conditions [123]. Unlike new batteries that can safely utilize a broad operational SoC range (e.g., 10%–90%), second-life systems are often constrained to tighter windows (e.g., 30%–70%) to limit stress and extend remaining useful life. This restriction directly reduces the effective usable energy fraction, thereby requiring oversizing in both energy (kWh) and power (kW) capacity relative to the predicted load demand of AI DCs. SLBESS must be sized not only for operational objectives such as FFR, peak shaving, and short-duration backup, but also to satisfy redundancy and fault-tolerance criteria. The SoH variability across SLBs introduces capacity dispersion and imbalance, which must be explicitly accounted for in system-level sizing methodologies [121].

Thermal and cycling considerations further complicate the sizing process. For SLBs with limited remaining useful life (RUL), constraining DoD and per-cycle energy throughput becomes essential to mitigate capacity fade and impedance growth. However, these operational safeguards further reduce usable capacity and necessitate additional oversizing. Moreover, increased internal resistance in aged lithium-ion cells leads to reduced power capability, particularly at lower SoC levels [124]. This reinforces the requirement for power oversizing to ensure that transient load spikes can be supported without violating voltage, thermal, or safety constraints.

From a deployment perspective, SLBESS are most suitable for the facility and grid layers of the proposed hierarchical energy storage architecture. At the facility level, SLBESS can be integrated with grid-interactive UPS systems to provide backup power, peak demand reduction, renewable energy smoothing, and short-term load buffering for AI data centers. At the grid level, larger-scale SLBESS installations can support ancillary services such as frequency regulation, reserve provision, congestion relief, and renewable energy integration, thereby improving both data center reliability and power system flexibility. Unlike chip- and rack-level storage, which require extremely high power density and ultra-fast response characteristics, SLBESS deployments prioritize energy capacity, economic value, and operational flexibility.

9. Challenges and future directions

This section synthesizes the challenges and future research directions arising from the multi-layer ESS integration reviewed in this paper. We first address current challenges for both AI DC operation and ESS integration, then discuss future technical directions. Table 5 provides a consolidated cross-technology comparison.

9.1. Current challenges for AI data centers

9.1.1. AI DC load characteristics and grid impact

As discussed in the Introduction and Section 2, AI DC load profiles differ fundamentally from traditional loads. The sub-second power variability of training workloads causes harmonic distortion, voltage flicker, and sub-frequency mechanical oscillations in synchronous generators [2,32]. These impacts are amplified when multiple GPU clusters train synchronously. At the transmission level, the sudden switch of large AI DCs to backup power, as demonstrated by the 1500 MW event in Virginia [10], reveals that current grid planning frameworks are insufficient for accommodating these highly responsive large loads. The primary challenge is the combination of high power magnitude, rapid and poorly predictable load transitions, and the absence of coordinated grid-response protocols between DC operators and utility grid operators.

9.1.2. Simulation and experimental validation

To validate the optimal operation of the energy management system (EMS) of an AI DC, it is necessary to model its different components, including the ESSs. To the best of the authors' knowledge, there is no software specifically developed for simulating EMSs of AI DCs. While MATLAB/Simulink or GT-SUITE can integrate mechanical and electrical components, the multi-physics nature of AI DCs, requiring simultaneous transient analysis for power electronics and thermal units alongside steady-state analysis for grid-scale systems, is not well supported by existing tools [33,125,126]. Available real-time simulator devices are similarly not suitable for real-time experimental analysis of AI DCs, since they are primarily developed for electrical circuits rather than combined electro-thermal-mechanical systems. This gap limits the ability to validate hierarchical ESS coordination strategies before physical deployment.

9.1.3. GPU scheduling and energy-aware compute

Current GPU-based AI workloads are predominantly executed at the highest supported core frequency, prioritizing peak performance at the expense of energy efficiency. Unlike CPUs, where DVFS is a mature and widely adopted mechanism, energy-aware frequency regulation for GPUs remains at an early stage [127,128]. Existing GPU scheduling policies largely optimize for throughput and latency without explicitly accounting for the energy consumption characteristics of heterogeneous deep learning tasks [129,130]. The resulting lack of controllability over GPU power profiles directly constrains the effectiveness of chip-level and rack-level ESSs in power smoothing applications, since these systems must cope with power profiles that are neither observable nor controllable from the ESS perspective.

9.1.4. Distribution transformers and bidirectional power flow

Traditional distribution transformers are designed for unidirectional power flow. AI DCs equipped with large-scale ESSs and advanced power electronics can dynamically shift between high consumption and active power injection for grid services. Under bidirectional power flow, conventional transformers experience accelerated thermal aging due to loading patterns that deviate from standard design assumptions [73]. Voltage regulation mechanisms and tap settings optimized for passive load behavior may become ineffective or unstable during rapid transitions between load and injection states [131]. These limitations underscore the need for enhanced transformer models, real-time monitoring, and power-electronic-aware grid interfaces at AI DC interconnection points.

Table 5
Comparative summary of ESSs solutions for AI DCs.

Technology	Response time	Application level	Key advantages	Key challenges	Relative cost	AI DC role	Grid interaction
Chip-level capacitors/HBM buffers	Sub-millisecond to millisecond	Component/Chip	Mitigates transient GPU power spikes and protects sensitive logic hardware.	Low energy storage capacity and localized thermal dissipation constraints.	Low per-unit cost; high aggregate cost.	Localized power smoothing for high-density processors.	None (isolated at the component level).
Rack BBU/supercapacitors	Millisecond to seconds	Rack/Server	Provides rapid transient response and supports modular, distributed deployment.	Limited energy capacity and accelerated degradation under high-frequency cycling.	Moderate.	Server-level power backup and rack-level load buffering.	None (isolated at the rack level).
Traditional UPS (TUPS)	<20 ms	Facility	Offers industry-proven reliability and complete critical load isolation.	Operates passively with idle battery assets; cannot provide grid ancillary services.	Moderate to High.	Primary load bridging during utility power interruptions.	Minimal power exchange with the external grid.
Grid-interactive UPS (GiUPS)	0.5–10 s	Facility	Enables bidirectional power flow and active grid support via fast frequency response.	Requires complex control systems and navigating regulatory approvals; increases battery degradation.	High.	Fast frequency response, voltage ride-through, and peak load shaving.	Active participation in frequency regulation and peak shaving markets.
Behind-the-Meter BESS (BTM BESS)	Minutes to hours	Facility/Data Center	Reduces peak demand charges and integrates on-site renewable energy sources.	High initial capital expenditure, large physical footprint, and complex capacity sizing.	High.	Total facility load smoothing, extended backup, and renewable energy integration.	Moderate interaction via the Point of Common Coupling (PCC).
Front-of-the-Meter BESS (FTM BESS)	Minutes to hours	Grid-scale	Provides massive energy capacity and wide-area ancillary grid services.	Subject to long permitting timelines, complex interconnection cues, and high capital costs.	Very High.	Utility-scale reserve capacity and long-term emergency power supply.	Continuous and strong bidirectional interaction with the transmission grid.
Fuel Cell (FC)	Seconds to minutes	Facility	Delivers low-carbon backup generation with high energy density to replace diesel generators.	Dependent on immature hydrogen supply chains and specialized infrastructure.	High.	Zero-emission standby generation and facility co-generation.	Moderate interaction depending on grid-parallel configuration.
Thermal Energy Storage (TES)	Minutes to hours	Facility	Enables shifting of cooling loads and utilizes off-peak electricity without degradation.	High integration complexity with existing HVAC systems and large physical space requirements.	Moderate.	Cooling system efficiency optimization and peak thermal load reduction.	Indirect interaction via shifted facility electrical demand.
Second-Life Battery ESS (SLBESS)	Seconds to hours	Facility/Data Center	Reduces initial capital costs by repurposing decommissioned electric vehicle batteries.	High uncertainty regarding State-of-Health tracking and lack of standardized safety metrics.	Low to Moderate.	Cost-effective behind-the-meter backup and supplemental energy storage.	Moderate interaction aligned with standard BESS operations.

9.1.5. Load forecasting for ESS sizing and dispatch

Traditional grid loads have predictable trajectories that can be forecast using historical data. AI DC load profiles are fundamentally different: training job submissions, checkpoint saves, and idle periods create highly stochastic power demand that cannot be reliably modeled with conventional forecasting techniques. As a result, ESS sizing methodologies based on historical load patterns may under- or over-estimate required capacity. Rack-level BESSs and chip-level batteries designed for power smoothing require accurate sub-second power consumption data, while grid-scale BESSs require accurate load profiles at the facility level for effective grid support. Advanced machine learning-based load forecasting methods [5,132–134] and high-resolution measurement infrastructure [1] are essential prerequisites for reliable ESS design and operation in AI DCs.

9.1.6. Advanced ESS degradation modeling and lifetime prediction

Due to high power variability in AI DCs, more frequent and irregular charging and discharging cycles for ESSs are inevitable. This requires precise degradation analysis to maximize lifetime and optimize ESS operation. As discussed in Section 7, ESSs at the chip or rack level experience high-frequency cycling under varying conditions to smooth the power profile. Online monitoring can support better lifetime prediction and proactive maintenance. For lithium-ion batteries, cyclic aging is a key concern at both the small scale (BBU, chip-level) and large scale (BESS). SLBESSs are particularly vulnerable to aging effects and require more rigorous monitoring, especially for larger units where cell-level heterogeneity amplifies uncertainty in state-of-health estimation [135, 136].

9.1.7. Hierarchical ESS coordination

Multi-layer ESS implementation in AI DCs requires a centralized EMS to monitor and coordinate operating modes across all ESS layers [137]. The on-site grid-scale BESS must operate coherently with the GiUPS to enable accurate power sharing, RES utilization, and emergency grid support. The GiUPS battery size and switching time are fundamentally coupled to the size of the on-site BESS. This relationship requires combined reliability and cost analyses. The work [138] demonstrates a novel control framework utilizing supercapacitors and grid-scale BESS for load smoothing, illustrating the potential of hierarchical coordination. TES systems must also operate coherently with the BESS and FC. BBUs and chip-level ESSs must coordinate for multi-layer power smoothing. These interdependencies necessitate a centralized control and management system with appropriate data acquisition infrastructure spanning all layers.

9.1.8. Long-duration energy storage (LDES)

On-site large-scale BESSs and GiUPS systems require long-duration energy storage (LDES) capability to compensate renewable intermittency or to sustain AI DC operation during extended utility outages [139]. Current LDES technologies are not yet mature for cost-effective deployment at the scale required by hyperscale AI DCs, and further investigation is needed into chemistries, system integration, and control strategies suitable for these applications.

9.1.9. Optimal sizing and cost considerations

There is limited public research on the optimal sizing and cost analysis of multi-layer ESSs in AI DCs. For instance, it is unclear whether a grid-scale BESS is more cost-effective than a GiUPS system with high power-density batteries, or whether large-scale supercapacitors at the rack level are preferable to a combination of BBUs and supercapacitors in load-smoothing scenarios. Answering these questions requires multi-objective techno-economic frameworks that integrate backup energy constraints, committed grid-service capacity, degradation-aware life-cycle cost modeling, and uncertainty in AI DC load profiles. Such frameworks are currently absent from the literature.

10. Future directions for AI data centers and energy storage

The future research needs for AI data centers (AI DCs) and energy storage systems (ESSs) can be organized into two main groups. The first group includes near-term directions that can support more practical deployment, monitoring, and operation of ESSs in AI DCs. The second group includes long-term directions that require deeper architectural changes, new control frameworks, and stronger coordination between AI DC operators, utilities, and market participants.

10.1. Near-term research needs

10.1.1. Power-aware GPU scheduling and ESS-coordinated operation

A fundamental near-term research direction is the co-design of GPU execution strategies with chip-level, rack-level, and facility-level ESSs. During training, it is essential to jointly optimize throughput and energy consumption through adaptive strategies for local and global batch size scheduling [140]. GPU resources should be dynamically scheduled across training tasks by accounting for time-varying electricity prices and renewable energy availability, enabling energy-aware model training [9,141–143].

More actionable research is needed on control algorithms that connect compute-side decisions, such as elastic GPU allocation and dynamic batch size selection, with ESS dispatch decisions. Elastic resource allocation, which dynamically adjusts the number of GPUs assigned to a training job, provides the controllability of GPU power consumption that is necessary for chip-level ESS or BBU scheduling. Beyond this, dynamic batch size selection can be implemented to regulate utility grid voltage profiles [144], directly linking compute scheduling to grid-level objectives. Therefore, an important near-term task is to develop real-time scheduling frameworks that jointly determine GPU allocation, batch size, and ESS charging/discharging actions under electricity price, renewable generation, and grid constraint signals.

10.1.2. AI DC load observability and forecasting infrastructure

Another near-term priority is the development of high-resolution observability infrastructure for AI DC loads. High-resolution monitoring of DC load dynamics requires waveform measurement units capable of capturing load profiles at very short time intervals, below 10 ms [1]. Such fine-grained measurements enable observation of fast transient behaviors that are invisible to conventional metering infrastructure.

Building on high-fidelity data, machine learning techniques can predict short-term and long-term load patterns to support proactive ESS control and optimization [5,132–134]. Advanced load modeling studies further increase the load observability of AI DCs for both short-term operation and long-term planning. A more actionable research path is to create standardized AI DC load datasets, define measurement requirements for sub-second and sub-cycle load variations, and develop forecasting models that can be directly used by energy management systems. Close collaboration between utilities and DC operators is also essential to enable secure, real-time access to load profiles, supporting coordinated grid–DC operation.

10.1.3. Advanced degradation modeling, state estimation, and RUL prediction

A further near-term research need is the development of degradation-aware ESS control methods for AI DC operating conditions. Lithium-ion batteries undergo both calendar and cyclic degradation [135]. Degradation pathways including SEI growth, lithium plating, and particle fracture [136] lead to capacity fade and increased internal resistance. For second-life batteries (SLBs) in particular, aging curves are shaped by first-life usage patterns, making each battery unique.

Estimating the remaining useful life (RUL) is essential for safe deployment, with approaches ranging from data-driven techniques [145] and neural network-based models to entropy-based algorithms [146]. A degradation-aware dispatch strategy is needed to avoid pushing cells into the non-linear knee-point region. Therefore, future work should focus on integrating state-of-health estimation, RUL prediction, and degradation cost directly into ESS dispatch optimization. The EU Battery Passport initiative [147] provides a lifecycle data framework that can inform second-life deployment decisions. Extending such frameworks to AI DC operating conditions, with their high cycling frequency and variable load profiles, is an important near-term research direction.

10.1.4. FTM grid-scale BESS as a reliable reserve asset

In the near term, more research is also needed on using front-of-the-meter (FTM) grid-scale BESSs at AI DC sites as reliable reserve assets. By monitoring and communicating the grid-scale BESS status at AI DC sites with grid operators, more optimized reserve planning can be achieved [148,149]. The total unused energy of the grid-scale BESS can be considered a potential reserve that can be injected into the grid during contingency events such as the tripping of synchronous generators. The FTM grid-scale BESS can also serve as an active reserve source in emergency cases by locally supplying the AI DC load in an islanded mode.

A more actionable research direction is to define dispatch protocols, communication requirements, and reserve qualification rules that allow AI DC-based BESSs to participate in grid services without compromising data center reliability. This requires precise coordination and control strategies between utility and AI DC operators, as well as appropriate communication and dispatch protocols. Research on market mechanisms and regulatory frameworks that allow AI DCs to participate in reserve markets remains limited.

10.2. Long-term research needs

10.2.1. Revisiting solid-state transformers for AI DC power delivery

A long-term research direction is to revisit solid-state transformers (SSTs) for high-density AI DC power delivery. SSTs [150–152] are primarily considered for enabling medium-voltage AC to high-voltage DC conversion, such as 800 V DC or higher, to improve rack-level power density in AI DCs [95]. In SST-based architectures, traditional UPS functionalities can be embedded within multi-stage power electronic converters, and the distribution transformer in the AI DC power delivery system can be eliminated, minimizing side effects of bidirectional power flow [8].

However, SSTs face challenges including lower reliability due to multi-level power conversion stages, limited power ratings, typically a few MVA, and increased control complexity and cost. While attractive for high-density power delivery, SSTs are not yet suitable replacements for UPS systems at gigawatt-scale DCs. Therefore, future long-term research should focus on modular SST architectures, reliability improvement, fault-tolerant converter control, protection design, and standardization. Multi-level modular SST architectures may improve scalability, but reliability and standardization remain key challenges requiring further investigation.

10.2.2. Integrated AI DC microgrid design and operation

A broader long-term direction is the design and operation of AI DCs as active microgrids. AI DC microgrids are emerging as a promising architecture for improving internal power resilience and external grid support [153,154]. These microgrids integrate BESS, conventional generators, distributed generation units, and utility grid connections, as illustrated in Fig. 9, to enhance reliability, redundancy, and operational flexibility. Within this framework, BESS placement and coordination are critical because properly located BESS can mitigate fast AI workload-induced load fluctuations, support local voltage and frequency stability, and coordinate the operation of heterogeneous generation resources.

From a topology perspective, two configurations are particularly relevant. In an *AC-bus* topology, GiUPS, BESS, on-site generation, and the

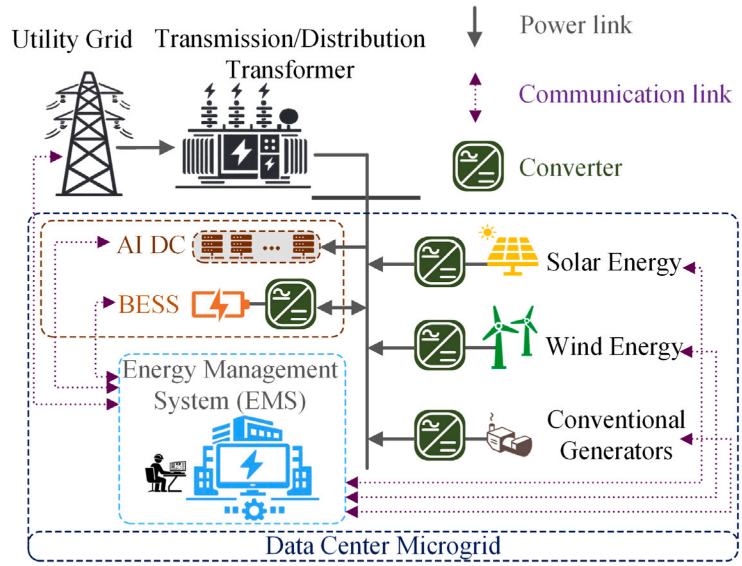

Fig. 9. AI DC microgrid with coordinated storage, generation, and energy management. The architecture integrates utility grid, BESS, renewable generation, conventional generation, and control links for flexible and resilient operation.

utility interconnection are coupled through a common medium-voltage AC bus, preserving compatibility with existing infrastructure but requiring tight synchronization among parallel inverter-based resources. In a *hybrid AC/DC* topology, a dedicated DC sub-bus directly supplies GPU rack loads, while the AC interface manages utility interaction and grid-service provision. This structure reduces unnecessary conversion stages, improves efficiency, and supports emerging 800 V DC rack architectures [8], but introduces additional protection-coordination challenges across the AC/DC boundary. Key practical issues include GFM–GiUPS coordination during grid-connected to islanded transitions and protection logic conflicts caused by GFM BESS fault-current injection. Initial mitigation strategies include master–slave control, virtual impedance, IEC 61850 GOOSE messaging, and adaptive relay settings supervised by the EMS [33,46].

Beyond internal reliability, BESS can enable AI DCs to provide grid services such as demand response, peak shaving, and fast reserves. This functionality blurs the traditional boundary between backup-oriented systems and active grid-support assets, particularly for GiUPS-based architectures. More broadly, the integration of GiUPS, grid-scale BESS, FC backup generation, TES for cooling, and renewable on-site generation under a unified hierarchical EMS suggests a future in which AI DCs operate as flexible grid participants rather than large disruptive loads. However, several challenges remain, including optimal microgrid topology design, hierarchical control under uncertainty, techno-economic co-optimization of ESS layers with compute scheduling, protection coordination for hybrid AC/DC systems, and regulatory frameworks for AI DC participation in grid services.

11. Conclusions

This paper has provided a comprehensive critical review of multi-layer ESSs for AI DCs and their roles in supporting both facility-level reliability and utility grid integration. Organized around a four-layer hierarchical taxonomy, namely chip-level buffers, rack/server-level ESSs, facility-level GiUPS systems, and grid-scale BESSs, the review systematically analyzes each technology's response timescale, application role, advantages, limitations, and coordination requirements.

Several key findings emerge from this review. First, AI DC load profiles are fundamentally different from traditional grid loads. Sub-second power variability, high peak-to-average ratios, and unpredictable workload transitions make conventional ESS dispatch strategies insufficient. Second, GiUPS systems represent the most significant architectural evolution in DC power infrastructure, transforming passive UPS batteries into active grid assets capable of FFR, VRT, peak shaving, and synthetic inertia provision. Third, grid-scale BESSs, whether BTM or FTM, are essential for demand charge reduction, renewable integration, and ancillary service provision at the facility level, but their sizing must explicitly account for degradation, load uncertainty, and multi-objective operational requirements. Fourth, non-battery technologies, namely FCs and TES, are underutilized in the current literature relative to their potential. FCs offer a viable clean replacement for diesel backup generation, while TES enables significant cooling efficiency gains and RES utilization. Fifth, SLBESSs offer a cost-effective deployment pathway for large-scale storage in AI DCs, provided that SoH uncertainty and cell heterogeneity are addressed through rigorous degradation-aware control.

The review identifies eight critical open challenges: the absence of multi-physics simulation tools for AI DC EMS validation, the immaturity of energy-aware GPU scheduling, bidirectional power flow impacts on distribution transformers, insufficient load forecasting methods for AI workloads, inadequate degradation models under AI DC cycling conditions, the lack of hierarchical ESS coordination frameworks, LDES technology immaturity, and the absence of multi-layer techno-economic sizing methodologies. Future research should address these challenges through the co-design of compute scheduling and ESS dispatch, the development of high-resolution load monitoring infrastructure, advanced degradation and RUL estimation methods for AI DC operating conditions, and integrated AI DC microgrid architectures. Addressing these open problems will be essential for transforming AI DCs from large disruptive grid loads into flexible, grid-supportive assets that contribute to power system reliability and the sustainable integration of renewable energy.

CRediT authorship contribution statement

Sina Mohammadi: Writing – review & editing, Writing – original draft, Visualization, Resources, Methodology, Investigation, Conceptualization. **Wayne Wang:** Writing – review & editing, Writing – original draft, Visualization, Resources, Methodology, Investigation. **Marcus Chen I Wada:** Writing – review & editing, Writing – original draft, Visualization, Resources, Methodology, Investigation. **Rouzbeh Haghighi:** Writing – review & editing, Writing – original draft, Visualization, Resources, Methodology, Investigation. **Ali Hassan:** Writing – review & editing, Writing – original draft, Visualization, Resources, Methodology, Investigation. **Hualong Liu:** Writing – review & editing, Writing – original draft, Visualization, Resources, Investigation.

Archit Bhatnagar: Writing – review & editing, Writing – original draft, Resources, Investigation. **Ang Chen:** Writing – review & editing, Supervision, Project administration. **Wencong Su:** Writing – review & editing, Supervision, Project administration.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

References

[1] Chalamala B, et al. Data center growth and grid readiness (TR131). Technical Report TR131, IEEE Power and Energy Society; 2025, http://dx.doi.org/10.17023/w4wy-s557.

[2] North American Electric Reliability Corporation (NERC). Characteristics and risks of emerging large loads. White Paper, NERC; 2025, URL: https://www.nerc.com/globalassets/who-we-are/standing-committees/rstc/whitepaper-characteristics-and-risks-of-emerging-large-loads.pdf. Large Loads Task Force (LLTF).

[3] Jonker A, Gomstyn A. What is an AI data center?. 2025, URL: https://www.ibm.com/think/topics/ai-data-center.

[4] Stream Data Centers. Data Center Development in an AI-Driven Market. Technical Report, Stream Data Centers; 2024, URL: https://www.streamdatacenters.com/wp-content/uploads/2024/02/SDC-BTPS-Whitepaper-240222.pdf.

[5] Lu S, Xiao C, Weng Y. Dynamic load model for data centers with pattern-consistent calibration. 2026, arXiv preprint arXiv:2602.07859.

[6] Vercellino R, Willard J, Campos G, Pereira WdS, Hull O, Selensky M, Mueller J. Measurement of generative AI workload power profiles for whole-facility data center infrastructure planning. 2026, arXiv preprint arXiv:2604.07345.

[7] Li X, Guan Y, Yao T, Wang Y, Wang W, Xu D. A soft-switching multiresonant switched-capacitor converter for data center applications. IEEE Trans Power Electron 2026;41(7):11079–97. http://dx.doi.org/10.1109/TPEL.2026.3657465.

[8] Xu J, Jiang X, Bao Y, Zheng Y, Chen X, Xu Q, Liao S, Ke D, Gao X. Sequential operating simulation of solid state transformer-driven next-generation 800 VDC data center. 2026, arXiv preprint arXiv:2601.16502.

[9] Chen X, Wang X, Colacelli A, Lee M, Xie L. Electricity demand and grid impacts of AI data centers: Challenges and prospects. 2025, arXiv preprint arXiv:2509.07218.

[10] McLaughlin T. Big tech's data center boom poses new risk to US grid operators. 2025, URL: https://www.reuters.com/technology/big-techs-data-center-boom-poses-new-risk-us-grid-operators-2025-03-19/.

[11] Ginzburg-Ganz E, Lifshits P, Machlev R, Belikov J, Krieger Z, Levron Y. Technical challenges of AI data center integration into power grids—A survey. Energies 2025;19(1):137.

[12] Rahman S, Khan TA. Energy storage systems for AI data centers: A review of technologies, characteristics, and applicability. Energies 2026;19(3):634.

[13] Safari A, Blaabjerg F, Oshnoei A. A research-industry perspective of battery systems technology for next-generation data centers. J Energy Storage 2026;152:120386. http://dx.doi.org/10.1016/j.est.2026.120386, URL: https://www.sciencedirect.com/science/article/pii/S2352152X26000502.

[14] Qian K, Xi Y, Cao J, Gao J, et al. Alibaba HPN: A data center network for large language model training. In: Proceedings of the ACM SIGCOMM 2024 conference. New York, NY, USA: Association for Computing Machinery; 2024, p. 691–706. http://dx.doi.org/10.1145/3651890.3672265, URL: https://dl.acm.org/doi/10.1145/3651890.3672265.

[15] Savant NV. GPU-to-GPU communication: Unlocking parallelism beyond the core. 2025, URL: https://medium.com/@nikheelvs/gpu-to-gpu-communication-unlocking-parallelism-beyond-the-core-a80de2974078.

[16] North American Electric Reliability Corporation (NERC). Characteristics and risks of emerging large loads. Technical Report, North American Electric Reliability Corporation; 2025, Large Loads Task Force White Paper.

[17] AtlanticNet. An overview of popular NVIDIA GPUs. 2025, URL: https://www.atlantic.net/gpu-server-hosting/an-overview-of-popular-nvidia-gpus/. [Accessed 27 February 2026].

[18] Rouslan Dimitrov NS, Blake M. How new GB300 NVL72 features provide steady power for AI. 2025, URL: https://developer.nvidia.com/blog/how-new-gb300-nvl72-features-provide-steady-power-for-ai/.

[19] Google Cloud. Accelerating AI inference with google cloud TPUs and GPUs. 2024, URL: https://cloud.google.com/blog/products/compute/accelerating-ai-inference-with-google-cloud-tpus-and-gpus. [Accessed 27 February 2026].

[20] Micron Technology. What changes in storage will AI drive?. 2024, URL: https://www.micron.com/about/blog/storage/ai/what-changes-in-storage-will-ai-drive.

[21] NVIDIA Corporation. GPUDirect storage: A direct path between storage and GPU memory. 2019, URL: https://developer.nvidia.com/blog/gpudirect-storage/.

[22] Gangidi A, Miao R, Zheng S, Bondu SJ, Goes G, Morsy H, Puri R, Riftadi M, Shetty AJ, Yang J, Zhang S, Fernandez MJ, Gandham S, Zeng H. RDMA over ethernet for distributed training at meta scale. In: Proceedings of the ACM SIGCOMM 2024 conference. ACM SIGCOMM '24, New York, NY, USA: Association for Computing Machinery; 2024, p. 57–70. http://dx.doi.org/10.1145/3651890.3672233.

[23] NVIDIA Corporation. Networking for data centers and the era of AI. 2023, URL: https://developer.nvidia.com/blog/networking-for-data-centers-and-the-era-of-ai/.

[24] Sheehan S, Rakow A. Evolving a data center into a microgrid: Industry perspectives and lessons learned. IEEE Electrification Mag 2023;11(3):16–25. http://dx.doi.org/10.1109/MELE.2023.3291193.

[25] US Department of Energy. Advantages and challenges of nuclear-powered data centers. 2025, URL: https://www.energy.gov/ne/articles/advantages-and-challenges-nuclear-powered-data-centers. [Accessed 27 February 2026].

[26] Nehrir H, Wang C. Hydrogen fuel and fuel cells: Potential candidates for sustainable, dispatchable, and environmentally friendly power generation and transport technologies. IEEE Energy Sustain Mag 2025;1(3):52–65. http://dx.doi.org/10.1109/ESM.2025.3606163.

[27] Vertiv. Understanding direct-to-chip cooling in HPC infrastructure: A deep dive into liquid cooling. 2024, URL: https://www.vertiv.com/en-us/about/news-and-insights/articles/educational-articles/understanding-direct-to-chip-cooling-in-hpc-infrastructure-a-deep-dive-into-liquid-cooling/.

[28] Takci MT, Qadrdan M, Summers J, Gustafsson J. Data centres as a source of flexibility for power systems. Energy Rep 2025;13:3661–71.

[29] Karpati A, Zsigmond G, Vörös M, Lendvay M. Uninterruptible power supplies (UPS) for data center. In: 2012 IEEE 10th jubilee international symposium on intelligent systems and informatics. IEEE; 2012, p. 351–5.

[30] Fawaz AH, Lorincz J, Mohammed AF. Minimizing data center uninterruptable power supply overload by server power capping. IEEE Commun Lett 2019;23(8):1342–6.

[31] 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 Convers Manage 2025;342:120137.

[32] Paananen J. Grid-Interactive Data Centers Enabling Energy Transition: Data center's hidden potential to provide essential grid services of a future power system. IEEE Electrification Mag 2023;11(3):26–34. http://dx.doi.org/10.1109/MELE.2023.3291195.

[33] Colangelo P, Coskun AK, Megrue J, Roberts C, Sengupta S, Sivaram V, Tiao E, Vijaykar A, Williams C, Wilson DC, et al. AI data centres as grid-interactive assets. Nat Energy 2025;1–8.

[34] Di Filippi A, Valentini L. How to Maximize Revenues from Your Data Center Energy Storage System with Grid Interactive UPS. Technical Report, Vertiv; 2025, Vertiv White Paper.

[35] Paananen J, Nasr E. Grid-interactive data centers: Enabling decarbonization and system stability. 2021, Dublin, Ireland.

[36] ONenergy. Hyperscale meets grid stability: ON.energy launches medium-voltage UPS built for AI data centers. 2025, URL: https://www.nacleanenergy.com/alternative-energies/hyperscale-meets-grid-stability-on-energy-launches-medium-voltage-ups-built-for-ai-data-centers.

[37] North American Electric Reliability Corporation. Fast Frequency Response Concepts and Bulk Power System Reliability Needs. Technical Report, North American Electric Reliability Corporation; 2020.

[38] Nordic Transmission System Operators. Technical Requirements for Fast Frequency Reserve Provision in the Nordic Synchronous Area. Technical Report, Statnett; 2021.

[39] Di Filippi A, Valentini L. How to maximize revenues from your data center energy storage system with grid interactive UPS. 2025, URL: https://www.vertiv.com/4918e5/globalassets/documents/white-papers/white-paper-maximize-revenues-data-center-energy-storage-grid-ups_329440_2.pdf.

[40] Eaton. Eaton and Microsoft's EnergyAware UPS technology pilot project. URL: https://www.eaton.com/us/en-us/products/backup-power-ups-surge-it-power-distribution/backup-power-ups/dual-purpose-ups-technology.html.

[41] Xie Y, Cui W, Wierman A. Enhancing data center low-voltage ride-through. 2025, arXiv preprint arXiv:2510.03867.

[42] Azizi A, Morovati S, Zamani A, Piruzza J, Guo D, Liu Z, Taimela P. Strengthening data center operations using grid-forming battery energy storage as a line-interactive uninterruptible power supply. Int J Electr Power Energy Syst 2026;175:111638.

[43] Zheng S, Teh J, Lai CM. Enhancing wind integration with battery energy storage systems reactive power support and dynamic line rating. IEEE Syst J 2026.

[44] Tozzi C. Voltage ride-through: A key ingredient in data center resilience. 2026, URL: https://www.datacenterknowledge.com/uptime/voltage-ride-through-a-key-ingredient-in-data-center-resilience.

[45] Conto J, Cheng Y, Rose J, Schmall J. Texas loads ride toward grid stability: Voltage ride through of large power electronic loads. IEEE Power Energy Mag 2025;23(5):56–67.
[46] Zahedi R, Zamani A, Anilkumar R. Best practices for large load interconnections: A North American perspective on data centers. 2026, arXiv preprint arXiv:2601.12686.
[47] Li R, Geng H, Yang G. Fault ride-through of renewable energy conversion systems during voltage recovery. J Mod Power Syst Clean Energy 2016;4(1):28–39.
[48] Cheng Y, Sahni M, Conto J, Huang SH, Schmall J. Voltage-profile-based approach for developing collection system aggregated models for wind generation resources for grid voltage ride-through studies. IET Renew Power Gener 2011;5(5):332–46.
[49] Pant D, Singh A, Van Bogaert G, Gallego YA, Diels L, Vanbroekhoven K. An introduction to the life cycle assessment (LCA) of bioelectrochemical systems (BES) for sustainable energy and product generation: relevance and key aspects. Renew Sustain Energy Rev 2011;15(2):1305–13.
[50] Kaiser R. Optimized battery-management system to improve storage lifetime in renewable energy systems. J Power Sources 2007;168(1):58–65.
[51] Zhang Y, Xu Y, Yang H, Dong ZY, Zhang R. Optimal whole-life-cycle planning of battery energy storage for multi-functional services in power systems. IEEE Trans Sustain Energy 2019;11(4):2077–86.
[52] Wang J, Ye S, Wu B, Liu B. Life-cycle performance analysis of a building integrated energy system considering equipment performance degradation. Energy Convers Manage 2026;347:120593.
[53] Kamali M, Hewage K. Life cycle performance of modular buildings: A critical review. Renew Sustain Energy Rev 2016;62:1171–83.
[54] Omar N, Monem MA, Firouz Y, Salminen J, Smekens J, Hegazy O, Gaulous H, Mulder G, Van den Bossche P, Coosemans T, et al. Lithium iron phosphate based battery–Assessment of the aging parameters and development of cycle life model. Appl Energy 2014;113:1575–85.
[55] Majeau-Bettez G, Hawkins TR, Strømman AH. Life cycle environmental assessment of lithium-ion and nickel metal hydride batteries for plug-in hybrid and battery electric vehicles. Environ Sci Technol 2011;45(10):4548–54.
[56] Ding Y, Cano ZP, Yu A, Lu J, Chen Z. Automotive li-ion batteries: Current status and future perspectives. Electrochem Energy Rev 2019;2(1):1–28.
[57] North American Electric Reliability Corporation (NERC). Grid forming functional specifications for BPS-connected battery energy storage systems. White Paper, North American Electric Reliability Corporation; 2023, URL: https://www.nerc.com/globalassets/our-work/reports/white-papers/white_paper_gfm_functional_specification.pdf.
[58] Donovan P. Understanding BESS: Battery Energy Storage Systems for Data Centers. Technical Report, (White Paper 185). Schneider Electric Energy Management Research Center; 2025.
[59] Azizi A, Aluko A, Maleki S, Salehi R, Bayat H. Grid-forming inverter applications for improved reliability and power quality in AI data centers. In: 2025 IEEE power & energy society general meeting. 2025, p. 1–5. http://dx.doi.org/10.1109/PESGM52009.2025.11225523.
[60] Du W, Tuffner FK, Schneider KP, Lasseter RH, Xie J, Chen Z, Bhattarai B. Modeling of grid-forming and grid-following inverters for dynamic simulation of large-scale distribution systems. IEEE Trans Power Deliv 2020;36(4):2035–45.
[61] Blaabjerg F, Teodorescu R, Liserre M, Timbus A. Overview of control and grid synchronization for distributed power generation systems. IEEE Trans Ind Electron 2006;53(5):1398–409. http://dx.doi.org/10.1109/TIE.2006.881997.
[62] Milano F, Dörfler F, Hug G, Hill DJ, Verbič G. Foundations and challenges of low-inertia systems (invited paper). In: 2018 power systems computation conference. 2018, p. 1–25. http://dx.doi.org/10.23919/PSCC.2018.8450880.
[63] Azizi A, Hooshyar A. Fault current limiting and grid code compliance for grid-forming inverters—Part I: Problem statement. IEEE Trans Sustain Energy 2024;15(4):2486–99.
[64] Lai CM, Teh J. Network topology optimisation based on dynamic thermal rating and battery storage systems for improved wind penetration and reliability. Appl Energy 2022;305:117837.
[65] Teh J, Lai CM. Reliability impacts of the dynamic thermal rating and battery energy storage systems on wind-integrated power networks. Sustain Energy Grids Netw. 2019;20:100268.
[66] Ahrabi RR, Mousavi A, Mohammadi E, Wu R, Chen AK. AI-Driven data center energy profile, power quality, sustainable sitting, and energy management: A comprehensive survey. In: 2025 IEEE conference on technologies for sustainability. IEEE; 2025, p. 1–8.
[67] Watson K. Data centers – A good grid citizen. 2025, Slide on grid support and batteries; Mission Critical Solutions presentation. Presentation.
[68] Quanta Tech LLC. Understanding AI load profiles and their impact on power systems. 2025, Online Seminar (Webinar).
[69] Yang L, Teh J, Lai CM. Assessing power grid vulnerability to extreme weather: The impact of electric vehicles and energy storage systems. J Energy Storage 2025;112:115534.
[70] Su Y, Teh J, Fang S, Dai Z, Wang P. Two-stage optimal dispatch framework of active distribution networks with hybrid energy storage systems via deep reinforcement learning and real-time feedback dispatch. J Energy Storage 2025;108:115169.
[71] Crusoe Energy, Redwood Materials. Crusoe and redwood materials power AI/Data center infrastructure using second-life EV battery microgrid. 2025, URL: https://www.crusoe.ai/resources/newsroom/crusoe-and-redwood-materials-power-the-future-of-ai.
[72] Tao X, Gadh R. Coordinated fast frequency response from electric vehicles, data centers, and battery energy storage systems. 2025, arXiv preprint arXiv:2512.14136.
[73] Majeed IB, Nwulu NI. Impact of reverse power flow on distributed transformers in a solar-photovoltaic-integrated low-voltage network. Energies 2022;15(23):9238.
[74] Pan L, Song M, Muzaffar N, Chen L, Ji C, Yao S, Xu J, Wu W, Li Y, Chen J, et al. Salt cavern redox flow battery: The next-generation long-duration, large-scale energy storage system. Curr Opin Electrochem 2025;49:101604.
[75] Mohamad F, Teh J, Lai CM. Optimum allocation of battery energy storage systems for power grid enhanced with solar energy. Energy 2021;223:120105.
[76] Yang L, Teh J, Alharbi B. Optimizing distributed generation and energy storage in distribution networks: Harnessing metaheuristic algorithms with dynamic thermal rating technology. J Energy Storage 2024;91:111989.
[77] Mamun A, Narayanan I, Wang D, Sivasubramaniam A, Fathy H. Multi-objective optimization of demand response in a datacenter with lithium-ion battery storage. J Energy Storage 2016;7:258–69.
[78] Ye G, Fang J, Wang N, Gaogao Y, Sun K. A grid-forming energy storage system capacity planning method considering device lifetime. Energies 2026.
[79] Yu Y, Shan K, Tang H, Wang S. Reliability and economic impacts of utilizing battery energy storage in data centers for energy flexibility services in smart grids. Energy Convers Manage 2025;339:119951.
[80] Cupelli L, Barve N, Monti A. Optimal sizing of data center battery energy storage system for provision of frequency containment reserve. In: IECON 2017 - 43rd annual conference of the IEEE industrial electronics society. 2017, p. 7185–90. http://dx.doi.org/10.1109/IECON.2017.8217257.
[81] Fan L, Tu Z, Chan SH. Recent development of hydrogen and fuel cell technologies: A review. Energy Rep 2021;7:8421–46.
[82] Manzo D, Thai R, Le HT, Venayagamoorthy GK. Fuel cell technology review: Types, economy, applications, and vehicle-to-grid scheme. Sustain Energy Technol Assess 2025;75:104229.
[83] Nikiforakis I, Mamalis S, Assanis D. Understanding solid oxide fuel cell hybridization: A critical review. Appl Energy 2025;377:124277.
[84] Kong IB, Kim WS, Chae S. A grid-forming control method for PEMFC power conversion systems with power ramp rate limitation to prevent fuel starvation. IEEE Open J Power Electron 2025.
[85] Perez F. Enabling net zero data centers: A techno-economic analysis of bloom energy's SOFC systems [Ph.D. thesis], Politecnico di Torino; 2025.
[86] Nikiforow K, Pennanen J, Ihonen J, Uski S, Koski P. Power ramp rate capabilities of a 5 kW proton exchange membrane fuel cell system with discrete ejector control. J Power Sources 2018;381:30–7.
[87] Li Y, Li Y. AI load dynamics–A power electronics perspective. 2025, arXiv preprint arXiv:2502.01647.
[88] Zheng P, Xie X, Zhang C, Cai S, Pan J, Zhang H, Yan M, Mu Q. Techno-economic assessment framework for 2.5 MW-scale grid-connected proton exchange membrane fuel cell power systems: A case study in China. Int J Hydrog Energy 2025;167:150942.
[89] Alva G, Lin Y, Fang G. An overview of thermal energy storage systems. Energy 2018;144:341–78.
[90] Enescu D, Chicco G, Porumb R, Seritan G. Thermal energy storage for grid applications: Current status and emerging trends. Energies 2020;13(2):340.
[91] Omrani M, Ghassemi M. AI-driven optimization of fan-wall cooling system in a medium-density data center. Int J Heat Mass Transfer 2025;247:127159.
[92] Guelpa E, Verda V. Thermal energy storage in district heating and cooling systems: A review. Appl Energy 2019;252:113474.
[93] Sarbu I, Sebarchievici C. A comprehensive review of thermal energy storage. Sustainability 2018;10(1):191.
[94] Luerssen C, Gandhi O, Reindl T, Sekhar C, Cheong D. Life cycle cost analysis (LCCA) of PV-powered cooling systems with thermal energy and battery storage for off-grid applications. Appl Energy 2020;273:115145.
[95] Huntington J, Tu M. 800 VDC architecture for next-generation AI infrastructure. 2025, URL: https://nvdam.nvidia.com/assets/share/asset/zlg5snufeo.
[96] IEC. Medium voltage DC (MVDC) grids for anall-electric society. 2025, URL: https://www.iec.ch/basecamp/medium-voltage-dc-mvdc-grids-all-electric-society.
[97] Sun D, Shapiro D, Kim B, Athavale J, Mercado R. Open compute project: Open rack V3 48v BBU (Rev 1.4). 2023, URL: https://www.opencompute.org/documents/open-rack-v3-bbu-module-spec-1-4-pdf.
[98] .Sun D, Shapiro D, Kim B, Athavale J, Mercado R. Open compute project: Open rack V3 BBU shelf (Rev 1.1). 2022, URL: https://www.opencompute.org/documents/open-rack-v3-bbu-shelf-spec-rev1-1-pdf-1.
[99] Dimitrov R, Petty H, Srivastava N, Blake M. How new GB300 NVL72 features provide steady power for AI. 2025, URL: https://developer.nvidia.com/blog/how-new-gb300-nvl72-features-provide-steady-power-for-ai.
[100] Analog Devices. ADI OCP ORV3 BBU reference design. 2024, URL: https://wiki.analog.com/resources/eval/adi-ocp-orv3-bbu-reference-design.

[101] Roach J. Microsoft datacenter batteries to support growth of renewables on the power grid. 2022, URL: https://news.microsoft.com/source/features/sustainability/ireland-wind-farm-datacenter-ups.
[102] Collath N, Tepe B, Englberger S, Jossen A, Hesse H. Aging aware operation of lithium-ion battery energy storage systems: A review. J Energy Storage 2022;55:105634. http://dx.doi.org/10.1016/j.est.2022.105634.
[103] Malla S, Deng Q, Ebrahimzadeh Z, Gasperetti J, Jain S, Kondety P, Ortiz T, Vieira D. Coordinated priority-aware charging of distributed batteries in oversubscribed data centers. In: 2020 53rd annual IEEE/ACM international symposium on microarchitecture. 2020, p. 839–51. http://dx.doi.org/10.1109/MICRO50266.2020.00073.
[104] Safari A, Sorouri H, Rahimi A, Oshnoei A. A systematic review of energy efficiency metrics for optimizing cloud data center operations and management. Electronics 2025;14(11):2214.
[105] Bjorlin A. OCP summit 2022: Open hardware for AI infrastructure (Grand teton / ORV3 rack and power). 2022, URL: https://engineering.fb.com/2022/10/18/open-source/ocp-summit-2022-grand-teton.
[106] Crawford AJ, Huang Q, Kintner-Meyer MC, Zhang JG, Reed DM, Sprenkle VL, Viswanathan VV, Choi D. Lifecycle comparison of selected li-ion battery chemistries under grid and electric vehicle duty cycle combinations. J Power Sources 2018;380:185–93. http://dx.doi.org/10.1016/j.jpowsour.2018.01.080.
[107] UL Solutions. UL 9540A Test Method for Battery Energy Storage Systems (BESS). URL: https://www.ul.com/services/ul-9540a-test-method.
[108] Choukse E, Warrier B, Heath S, Belmont L, Zhao A, Khan HA, Harry B, Kappel M, Hewett RJ, Datta K, et al. Power stabilization for AI training datacenters. 2025, arXiv preprint arXiv:2508.14318.
[109] Genkina D. Will Supercapacitors Come to AI's rescue? Power bursts in large AI workloads can threaten to overwhelm the grid. 2025, URL: https://spectrum.ieee.org/supercapacitor-2671883490.
[110] Eaton, b. Supercapacitors in AI Data Centers. URL: https://www.eaton.com/us/en-us/products/electronic-components/infographics/supercaps-in-ai-datacenters.html.
[111] NVIDIA. NVML API reference guide - GPU deployment and management documentation. 2025, URL: https://docs.nvidia.com/deploy/nvml-api/group__nvmlDeviceQueries.html#group__nvmlDeviceQueries_1gf754f109beca3a4a8c8c1cd650d7d66c.
[112] Kaur J, Bath SK. Harmonic distortion in power systems due to electronic control and renewable energy integration: a comprehensive review. Discov Electron 2025;2(1):67. http://dx.doi.org/10.1007/s44291-025-00111-9.
[113] UL. UL 1974: Evaluation for repurposing batteries. 2024, URL: https://www.ul.com/services/second-life-electric-vehicle-battery-repurposing-facility-certification.
[114] Wang H, Rasheed M, Hassan R, Kamel M, Tong S, Zane R. Life-extended active battery control for energy storage using electric vehicle retired batteries. IEEE Trans Power Electron 2023;38(6):6801–5.
[115] White C, Thompson B, Swan LG. Repurposed electric vehicle battery performance in second-life electricity grid frequency regulation service. J Energy Storage 2020;28:101278. http://dx.doi.org/10.1016/j.est.2020.101278.
[116] Norris TH, Profeta TH, Patino-Echeverri D, Cowie-Haskell A. Rethinking load growth: Assessing the potential for integration of large flexible loads in US power systems. Report, Durham, NC: Nicholas Institute for Energy, Environment & Sustainability, Duke University; 2025, URL: https://hdl.handle.net/10161/32077. Report NI R 25-01.
[117] Roy B. New power players: How big tech firms are disrupting energy markets. 2025, URL: https://www.uh.edu/energy/news/stories/2025/ow_big_tech_firms_are_disrupting_energy_hmarkets.php.[Accessed 24 January 2026].
[118] BloombergNEF. Lithium-ion battery pack prices fall to $108 per kilowatt-hour despite rising metal prices. 2025, URL: https://about.bnef.com/insights/clean-transport/lithium-ion-battery-pack-prices-fall-to-108-per-kilowatt-hour-despite-rising-metal-prices-bloombergnef/.
[119] Bach A, Onori S, Reichelstein S, Zhuang J. Fair market value of used capacity assets: Forecasts for repurposed electric vehicle batteries. Technical Report, ZEW – Centre for European Economic Research; 2025, Discussion Paper No. 25-065.
[120] Internal Revenue Service (IRS). Domestic content bonus credit. 2025, URL: https://www.irs.gov/credits-deductions/clean-electricity-production-credit.
[121] Casals LC, García BA, Canal C. Second life batteries lifespan: Rest of useful life and environmental analysis. J Environ Manag 2019;232:354–63.
[122] Martinez-Laserna E, Gandiaga I, Sarasketa-Zabala E, Badeda J, Stroe DI, Swierczynski M, Goikoetxea A. Battery second life: Hype, hope or reality? A critical review of the state of the art. Renew Sustain Energy Rev 2018;93:701–18.
[123] Neubauer J, Pesaran A. The ability of battery second use strategies to impact plug-in electric vehicle prices and serve utility energy storage applications. J Power Sources 2011;196(23):10351–8.
[124] Vetter J, Novák P, Wagner MR, Veit C, Möller KC, Besenhard J, Winter M, Wohlfahrt-Mehrens M, Vogler C, Hammouche A. Ageing mechanisms in lithium-ion batteries. J Power Sources 2005;147(1–2):269–81.
[125] Cao Z, Li M, Lin F, Jia J, Wen Y, Yin J, See S. Transforming future data center operations and management via physical AI. 2025, arXiv preprint arXiv:2504.04982.
[126] Li M, Wang R, Tan R, Wen Y. Phythesis: Physics-guided evolutionary scene synthesis for energy-efficient data center design via LLMs. 2025, arXiv preprint arXiv:2512.10611.
[127] Chung JW, Wu R, Ma JJ, Chowdhury M. Where do the joules go? Diagnosing inference energy consumption. 2026, arXiv preprint arXiv:2601.22076.
[128] Wu R, Chung JW, Chowdhury M. Kareus: Joint reduction of dynamic and static energy in large model training. 2026, arXiv preprint arXiv:2601.17654.
[129] Chung JW, Gu Y, Jang I, Meng L, Bansal N, Chowdhury M. Reducing energy bloat in large model training. In: Proceedings of the ACM SIGOPS 30th symposium on operating systems principles. 2024, p. 144–59.
[130] Chung J-W, Ma JJ, Wu R, Liu J, Kweon OJ, Xia Y, Wu Z, Chowdhury M. The ML. ENERGY benchmark: Toward automated inference energy measurement and optimization. 2025, arXiv preprint arXiv:2505.06371.
[131] Frotscher R, Rave M, teNyenhuis E, Upadhyay P. Reverse power flow impact on transformers. 2021, URL: https://grouper.ieee.org/groups/transformers/subcommittees/standardsc/C57.133/F24-C57.133-ReversePowerFlowTutorial-%20teNyenhuis.pdf. Slides. IEEE PES Transformers Committee Spring 2021 Meeting – Tutorial / Technical Presentation.
[132] Mughees M, Li Y, Chen Y, Li YR. Short-term load forecasting for AI-data center. In: 2025 IEEE power & energy society general meeting. IEEE; 2025, p. 1–5.
[133] Jiang H, Qu B, Zhu J, Zeng F, Lin X, Zhong W. HyperLoad: A cross-modality enhanced large language model-based framework for green data center cooling load prediction. 2025, arXiv preprint arXiv:2512.19114.
[134] Mo R, Lin W, Liu G, Liu H, He L. Learning from imbalance: Cross-server power prediction in large data centers via domain adaptation regression. Expert Syst Appl 2025;287:127845.
[135] Kostenko GP. Accounting calendar and cyclic ageing factors in diagnostic and prognostic models of second-life EV batteries. 2024.
[136] Hassan A, Khan SA, Li R, Su W, Zhou X, Wang M, Wang B. Second-life batteries: A review on power grid applications, degradation mechanisms, and power electronics interface architectures. Batteries 2023;9(12):571.
[137] Li L, Yang P, Ju Y, Hu Z, Wang Z. Coordinating multiple UPS batteries in datacenter for load flexibility and cost reduction. SSRN Electron J 2025. http://dx.doi.org/10.2139/ssrn.5647230, URL: https://ssrn.com/abstract=5647230.
[138] Ko MS, Shim JW, Zhu H. Mitigation of datacenter demand ramping and fluctuation using hybrid ESS and supercapacitor. 2025, arXiv preprint arXiv:2512.08076.
[139] Farzan F, Mahani K, Farzan F, Masiello R, Brown W. Long-duration energy storage: Planning and operation to enhance power grid sustainability. IEEE Energy Sustain Mag 2025;1(3):41–51. http://dx.doi.org/10.1109/ESM.2025.3606162.
[140] Gu D, Xie X, Huang G, Jin X, Liu X. Energy-efficient GPU clusters scheduling for deep learning. 2023, arXiv preprint arXiv:2304.06381. arXiv:2304.06381. URL: https://arxiv.org/abs/2304.06381.
[141] Wang Y, Guo Q, Chen M. Providing load flexibility by reshaping power profiles of large language model workloads. Adv Appl Energy 2025;100232.
[142] Kez DA, Foley A. Instability risks from programmable AI load ramping in low-inertia grids. SSRN Electron J 2025. http://dx.doi.org/10.2139/ssrn.5370875, URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5370875.
[143] Crozier C, Liska M. The potential of data center energy demand to provide grid flexibility. Curr Sustainable/Renewable Energy Rep 2025;12(1):12.
[144] Liang Z, Chung JW, Chowdhury M, Chen J, Dvorkin V. GPU-to-grid: Voltage regulation via GPU utilization control. 2026, arXiv preprint arXiv:2602.05116.
[145] Cui X, Khan MA, Singh S, Sharma R, Onori S. Toward a BMS 2.0 design framework: Adaptive data-driven state-of-health estimation for second-life batteries with BIBO stability guarantees. IEEE Trans Transp Electrification 2025.
[146] Strugnell-Lees B, Evdokimova E, Wik T. An entropy-based, self-adaptive predictive algorithm for battery degradation. 2024.
[147] European Union. Regulation (EU) 2023/1542. 2023, URL: https://eur-lex.europa.eu/eli/reg/2023/1542/oj.
[148] Xu B, Wang Y, Dvorkin Y, Fernández-Blanco R, Silva-Monroy CA, Watson JP, Kirschen DS. Scalable planning for energy storage in energy and reserve markets. IEEE Trans Power Syst 2017;32(6):4515–27. http://dx.doi.org/10.1109/TPWRS.2017.2682790.
[149] Padmanabhan N, Ahmed M, Bhattacharya K. Battery energy storage systems in energy and reserve markets. IEEE Trans Power Syst 2020;35(1):215–26. http://dx.doi.org/10.1109/TPWRS.2019.2936131.
[150] Rothmund D, Guillod T, Bortis D, Kolar JW. 99% efficient 10 kV SiC-based 7 kV/400 V DC transformer for future data centers. IEEE J Emerg Sel Top Power Electron 2019;7(2):753–67. http://dx.doi.org/10.1109/JESTPE.2018.2886139.
[151] Sidorov A, Zinoviev G, Petzoldt J. Solid state transformer as a part electrical circuit for data center application. In: 2021 IEEE 22nd international conference of Young professionals in electron devices and materials. IEEE; 2021, p. 354–9.
[152] Huang AQ. Medium-voltage solid-state transformer: Technology for a smarter and resilient grid. IEEE Ind Electron Mag 2016;10(3):29–42.
[153] Irion J, Wiesner P, Bader J, Kao O. Optimizing microgrid composition for sustainable data centers. In: Proceedings of the SC'25 workshops of the international conference for high performance computing, networking, storage and analysis. 2025, p. 1990–6.
[154] 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.