

The evolving role of energy storage systems in AI data center load management

Part two of a two-part series

By Seth Miller, Ben Campbell, Joshua Sorensen

March 12, 2026

Key takeaways

- Lithium-ion (Li-ion) remains the main battery chemistry for battery energy storage systems (BESS), but sodium-ion (Na-ion) and flow batteries could soon see adoption as backup power storage.
- Data centers can use software to distribute compute tasks across multiple sites and match workloads to available power.
- Rapid improvements in processing architectures and algorithms, and the value proposition of BESS and other back-up systems, will evolve against this highly volatile backdrop.
- Stricter data center-specific tariffs create financial risks, so operators need to increase investment in behind-the-meter power generation and BESS.

Processing workload is a commodity. It's a digital product that leaves no trace of what systems produced it or where it was created. Commodity economics suggest that the data center industry should adopt standards that expand procurement, improve reliability, and reduce overall risk.

Standardization will happen over time. For now, hyperscalers and their suppliers are scrambling to put projects on the grid and design more efficient, scalable solutions. No single power architecture can address every challenge. Instead, operators must evaluate a mix of solutions like:

- Batteries, generators, and supercapacitors

- Software-based controls
- Using DC electrical architecture instead of AC

Each technology offers distinct advantages at different time scales and operating conditions. Data center power has long been a major part of the energy-storage market and has unique performance requirements compared to other applications. In the coming years, we expect a wider range of solutions instead of consolidation through standardization, as the industry responds to new opportunities and challenges. The industry may spend hundreds of billions of dollars on infrastructure each year—perhaps exceeding a trillion dollars annually by the end of the decade—it's still far from mature.

In the near term, the energy-storage market will likely stay segmented. Vendors will align their products with specific GPU providers like Nvidia, AMD, Google, and Amazon, and with hyperscalers that design the data center power and cooling systems. The most likely outcome is a segmented market dominated by larger energy-storage system developers that can match their development cycles to fast-moving clients and deploy reliable systems at scale. Even so, the industry still seek niche solutions that could compete on total cost of ownership (TCO), performance, and long-term reliability.

Read part one of this series

For more information about data centers, read part one of this series [The evolving role of energy storage systems in AI data center load management](#). In part one, we discuss how power dynamics limit data center growth and the four major power challenges that data centers face.

Competing solutions and market risks

Li-ion batteries

Li-ion batteries remain a foundational technology for data center uninterruptible power supply (UPS) systems and grid-forming BESS. Li-ion batteries are widely used because they offer:

- High power density
- Fast inverter-driven response at about 30 milliseconds (ms)
- A mature, safety-certified supply chain

Li-ion's low TCO, smaller footprint, and long life cycle have let it gradually replace lead-acid in UPS blocks. Lithium iron phosphate (LFP) is the main Li-ion chemistry for data centers because it's safer than other Li-ion systems. But it's still flammable and requires significant compromises in data center design. Manufacturers are tailoring LFP products for data center applications. For example, tier-one suppliers like LG are developing LFP cells that enable a 12C (five-minute) discharge rate. LG reports that these designs halve the number of cabinets operators need to serve a 1-megawatt (MW) UPS load and that [they're capable of delivering up to](#)

[20C of power output](#) in bursts.

The physical integration of power into next-generation data centers is in flux. Today, operators typically place UPS systems in a separate room from the data hall where the GPU racks are located, but these layouts may change as battery and GPU technologies evolve.

Battery backups (BBUs) are collocated with the racks. Locating storage and chips close together can reduce UPS costs and improve reliability by reducing the “blast radius” of a storage failure to a single rack rather than an entire hall. But this approach adds cost because the local batteries require fire suppression units installed at the racks themselves. Space is also a challenge because the copper cables connecting chips can extend only about 2 meters before signal propagation delays interrupt real-time communication. So BBUs can only be placed above or below a rack of chips, and some designs can’t use them at all.

The most common data center designs separate UPS systems from processing hardware because of physical space requirements. But operators might reconsider this layout as advances in optical interconnects enable real-time communication between chips at greater distances or as denser battery systems become available. Considerations like fire suppression or processing workload density may push data centers to continue to rely on separate UPS for LFP batteries, but some capacitors or supercapacitors will be present at the rack in any architecture.

Operators could shift some of the backup and load-smoothing needs from UPS systems to larger-scale BESS installations that are physically separate from the data center, but still capable of supplying power quickly. NERC’s [Welcome to the Large Loads Task Force Meeting and Workshop!](#) (PDF) presentation mentions Tesla’s proposal to equip data centers with grid-forming BESS with response times of 30 ms. These systems can’t fully meet the 10-ms UPS response time but are close enough to complement it. With such a configuration, operators could downsize the more expensive UPS substantially.

Rapid improvements in GPU performance and IT infrastructure are shaping predictions for future data center design. In a [discussion with Dwarkesh Patel](#), Microsoft CEO Satya Nadella cited the need to continually reassess assumptions about capital infrastructure as a reason Microsoft slowed the pace of its planned AI expansion. In this discussion, Nadella mentioned rapid changes in scale, power needs for racks and rows, and cooling requirements. Future data centers may require an entirely different physical layout than what’s optimal today. While the physical data center depreciates more slowly than the GPUs inside, it’s not clear whether the power, cooling, or even walls in today’s data centers will need to be torn out after this generation of GPUs is retired.

Motivations and barriers for alternative technology adoption

Li-ion still needs significant fire suppression and thermal management. Ordinary BESS based on LFP technology must sit idle to cool between charge and discharge cycles, and their heat output limits how quickly

they can operate. New storage technologies may benefit AI data centers if operators can install them economically at scale. But data center owners have been highly risk-averse. They must weigh the benefits of peak shaving against the risk of relying on emerging technology for perform critical reliability needs.

Simple diurnal peak-shifting can give new technologies a chance to prove themselves because the cost of failure is simply a higher demand charge. If the technology is stable, operators can integrate its capabilities into the data center's power architecture. But this process will be slow because of the huge revenue streams that flow downstream from power. We see hyperscalers typically making the decisions on these investments. The architectural considerations vary by site. We expect it'll take time to standardize technology choices and system designs since technology is developing so quickly. Still, new storage technologies will appear in data centers as pilot projects soon.

Solid-state, lithium-metal batteries (LiM-SSBs)

LiM-SSBs may gain a foothold in the market by offering:

- Better safety margins
- Lower thermal-runaway risk
- Potentially higher energy density

These advantages could make LiM-SSBs appealing in dense AI facilities where UPS footprint and fire codes are binding constraints.

The biggest challenge for introducing any new storage technology, including LiM-SSBs, to data centers is the tradeoff between maturity and safety.

The biggest challenge for introducing any new storage technology, including LiM-SSBs, to data centers is the tradeoff between maturity and safety. This is similar to when Li-ion systems first tried to displace lead-acid batteries in data centers at the end of the 2010s. Li-ion systems were superior for energy density, power density, and cycle life. But operators were hesitant because Li-ion systems were unfamiliar, expensive, and operationally unproven. A [2018 white paper from Vertiv](#) illustrates how those limited safety and reliability metrics slowed the initial adoption of Li-ion batteries for use in data centers. Industry safety and reliability requirements are even stricter today, as the value of capital-stack battery support has risen.

Even with these challenges, LiM-SSBs may gain a foothold in the market, because:

- They may appeal to operators who want to diversify the supply chain and reduce long-term procurement risks, because they reduce exposure to China and Foreign Entity of Concern restrictions.
- Higher thermal tolerance and energy density enable denser deployments at the rack.

- Manufacturers can adapt LiM-SSBs they designed for EVs for use in data centers, where high operating temperatures may improve solid-state electrolyte performance.
- Reducing fire-suppression infrastructure could tip data centers from centralized UPS rooms back to BBUs, if inherently safe batteries become available.
- Potentially higher C-rates and pulse-power capabilities could remove the need for supercapacitors for UPS and fast-switching applications.
- Higher voltages may reduce the number of cells needed for high-voltage bus architecture and improve reliability.

Sodium-ion (Na-ion)

Na-ion is emerging as a potentially safer and lower-cost alternative to Li-ion. Na-ion has a competitive long-term cost structure and supply chain to compared to Li-ion, which is highly sensitive to the cost of lithium. Some Na-ion chemistries also have better thermal stability than LFP, which may reduce fire-suppression needs—an important consideration for data centers with high rack densities and tight permitting conditions. Some Na-ion chemistries exhibit good reliability and cycle life even under high temperature cycling, which allows for more daily cycles than Li-ion (though not as many as flow batteries). Na-ion offers a balanced middle ground among energy density, safety, and cycling performance.

Na-ion offers a balanced middle ground among energy density, safety, and cycling performance.

Na-ion batteries also claim higher power performance than Li-ion batteries. The now-defunct startup Natron received attention in the data center space for its Na-ion cells with higher power and excellent cycle life. Natron offered a unique Na-ion chemistry for supercapacitor-like performance. The chemistry required unique handling, and Natron was unable to scale their production.

Na-ion has limitations. It's early in commercialization and has lower energy density than Li-ion. But as Na-ion proves itself in utility markets, data center operators may see those installations as proof the technology is ready for data centers. Na-ion is well-suited for:

- High-power UPS applications
- Cost-optimized peak shaving
- Moderate load shaping
- Future interconnection-flexibility uses
- Integration in BBUs, especially if safety improves further

We expect an increasing amount of Chinese supply to ramp up to target this application.

Flow batteries

Flow batteries have several features that fit the direction of data center power architectures. Vanadium and

other aqueous flow chemistries use nonflammable electrolytes. This avoids the fire-propagation risks that limit placement of Li-based systems in or near high-density computer halls. As GPU clusters become more compact to support high-bandwidth interconnections, operators are removing BBUs from the data room. Flow batteries avoid limitations that constrain Li-ion, like:

- Thermal-load limits
- NFPA-driven separation requirements
- Class C mitigation systems

These advantages makes it easier to site flow batteries near centralized UPS rooms.

In this long-duration UPS structure, flow batteries would serve as a stable, long-duration backbone that complements high-power Li-ion applications while reducing the need for longer-duration Li-ion systems.

Although flow batteries aren't designed for millisecond-scale ride-through, validation work from TerraFlow Energy, described in its paper [Fast Response Validation of Storing Flow Battery Stacks for LDUPS Systems](#) (PDF), shows that modern stacks can respond rapidly enough to handle the slower parts of UPS operations—from several seconds to several minutes. It also found that modern stacks can cycle repeatedly without the degradation Li-ion batteries experience under similar conditions. Decoupling power and energy lets operators do two things:

- Size a vanadium redox flow power stack for the desired response window
- Independently scale electrolyte tanks (which define energy storage size) for throughput

This enables a nondegrading buffer for load shaping, charge management, and interconnection-related duty cycles that AI campuses need.

As data centers evolve, flow batteries may also create a long-duration UPS opportunity. Many operators are moving back to centralized UPS systems because denser, thermally contained GPU clusters can't accommodate BBU modules. In a layered UPS system, Li-ion or ultracapacitor modules provide the subsecond ride-through while a flow battery—validated to sustain high power for extended windows—takes over for minutes to hours. The long-duration storage made possible by this architecture would:

- Reduce diesel generator runtime
- Mitigate fuel constraints
- Improve resilience without the safety risks that come with scaling Li-ion

In this long-duration UPS structure, flow batteries would serve as a stable, long-duration backbone that complements high-power Li-ion applications while reducing the need for longer-duration Li-ion systems.

Supercapacitors and Li-ion capacitors

Purely electronic capacitors are the fastest component of a UPS system, responding in less than a millisecond. Supercapacitors store charge through the motion of ions in an electrolyte, like a battery, but accept ions at the electrode surface rather than its bulk. This eliminates kinetic barriers to ion insertion into the electrode that slow movement. As a result, supercapacitors respond in milliseconds and offer volumetric energy density about two orders of magnitude higher than electronic capacitors. Operators are scaling supercapacitor banks into the hundreds of MW to stabilize data center frequencies. The best supercapacitors have an energy density that's about 10 watt-hours per liter (Wh per L) and offer 10 to 20 years of cycle life.

Li-ion capacitors (LiCs) offer step up in energy density. A LiC uses a high-surface-area cathode like a supercapacitor but substitutes a lithium-containing electrolyte and graphite anode like those in standard Li-ion batteries. These systems offer similar power capability to supercapacitors, but with a higher energy density of 30 to 50 Wh per L. We believe the TCO of a LiC is higher than that of a supercapacitor, but details aren't readily available and the value from higher energy density may support the increased cost. Vendors like [Delta Power](#) are shipping LiCs that can sustain power for 5 to 15 seconds as a bridge to backup power in data centers.

The European company Skeleton Technologies is producing a SuperBattery capable of charging at 40°C. The battery appears to be a LiC with a [niobium oxide-containing anode](#). The company reports up to [48 Wh per L](#) in energy density and a lifetime over 50,000 cycles. This density is about a third of conventional BESS offerings from companies [like the Chinese company CATL](#), which advertised for the latest generation of LFP cells for BESS. But Skeleton's cells, with their extended lifetime, will last much longer than Li-ion cells, which are likely to survive only a few thousand cycles when operated at similar temperatures.

Check out our tools for more information

The E Source [Battery Forecast Database](#) provides visibility in our battery demand and battery price forecasts for over 40 battery applications, including data on 12 stationary applications.

The E Source [Battery Factory Tracker](#) is an analysis of all active and announced Li-ion battery factories. You can use the data to improve your understanding of the manufacturing landscape.

Nonstorage solutions and risks

Besides new energy-storage technologies, data centers have other solutions addressing their four core power problems. Among these nonstorage solutions are software platforms, tools that manage interconnection bottlenecks, and more-advanced AI chips.

High Voltage DC architectures

Data centers have historically distributed electricity inside data center server racks using 12-volt (V) power supplies, an architecture inherited from the desktop PC. This worked when racks consumed 5 to 15 kilowatts (kW). As processors grew more power-hungry—particularly multi-core CPUs and early GPU accelerators—the currents required at 12V became punishing. Google recognized this problem as early as 2009, began using 48V DC infrastructure. By 2016, Google had joined the Open Compute Project (OCP) to standardize a 48V rack architecture. Google reported a 30% reduction in conversion losses and a 16-fold reduction in distribution losses compared to 12V. Today, 48V (or 54V in NVIDIA's convention) is the standard for hyperscale operators. The OCP Open Rack v3 specification codifies 48V power shelves, rack-level Li-ion BBUs, and busbar distribution as the baseline architecture for modern data centers.

The same limitations that drove the 12V-to-48V transition is now forcing a move to even higher voltages. AI training racks are pushing power envelopes to 100–600 kW today and may grow to 1 MW per rack by 2027. At 48V, delivering 600 kW requires approximately 12,500 amps—a current level that demands busbars weighing close to 200 pounds and may need liquid cooling just for the power conductors. The proposed solution is to move to $\pm 400\text{V}$ (effectively 800V pole-to-pole) high voltage DC (HVDC) distribution, where that same 600 kW requires only about 750 amps and allows for air-cooled busbars with an 85% weight reduction.

Beyond the copper savings, HVDC distribution reduces the power conversion stages. In traditional AC architecture, power is rectified to DC in the UPS, inverted back to AC, stepped down through a transformer in the power distribution unit, then re-rectified in each server's power supply. The transition between each stage loses 1–3% efficiency. A 400V DC bus eliminates the UPS inverter stage, the power distribution unit transformer, and the server's front-end rectifier and power factor correction circuit, yielding an overall 7–10% efficiency gain. The industry is converging rapidly:

- Microsoft and Meta launched the “Mt. Diablo” OCP specification in late 2024 (and Google joined in 2025), which standardized a $\pm 400\text{V}$ three-conductor system as a transitional architecture.
- NVIDIA's Kyber rack platform targets native 800V DC distribution with production deployments in 2027.
- Power electronics vendors like Vertiv, Delta, and Eaton have all announced 800V DC product lines for 2026

The shift to HVDC distribution changes how batteries fit into the data center power chain—and arguably makes them more important. In the traditional AC architecture, batteries sit inside a central UPS that performs a wasteful double conversion: AC is rectified to DC to charge the batteries, then the battery output is inverted back to AC for distribution. In a 400V DC architecture, a BBU connects directly to the DC distribution bus with no inverter or rectifier, since a Li-ion battery string (roughly 100–110 series cells of LFP chemistry) naturally produces a voltage in the 320–400V range. These batteries further protect the system from large voltage swings when the chips switch operational modes. Also, a native DC bus simplifies integration with on-site solar, fuel cells, and grid-scale BESS, which all produce DC natively and can feed the bus without additional

inversion stages. The HVDC transition, in short, doesn't just accommodate batteries—it architecturally privileges them. Because these batteries will be integral to the function of the data center, vendors are likely to lean toward the most reliable, field-tested chemistries, favoring LFP. But the advantages of collocating BBUs with racks may present an opportunity for inherently nonflammable chemistries as they prove themselves over the next decade.

Software solutions for load flexibility

An AI data center produces information, so data centers can rapidly switch loads between locations to balance temporary grid strain in one location with availability in another. Analysis by the Electric Power Research Institute in its 2024 white paper [Powering Intelligence: Analyzing Artificial Intelligence and Data Center Energy Consumption](#) (PDF) suggests that widely deployed data centers could stabilize the grid by responding to electricity supply conditions. Software that modulates processing workload in response to grid signals partly enables this stability. This software creates a flexible demand resource that can complement battery storage in congestion and balancing applications

Software platforms like Emerald AI's [Emerald Conductor Platform](#) will move or delay calculations to prioritize uninterruptible tasks. These tasks, like commercial training, run over relocatable jobs like inference and pause the least time-sensitive results like academic research. Startup [Hammerhead AI](#) makes data centers more efficient by predicting how much power is available in the grid for processing workload. These systems use software orchestration instead of hardware to flatten shallow changes in power demand that would otherwise need a battery buffer. The primary goal is the same for each:

- Let data centers stand up or expand operations in regions with unreliable supply
- Improve the return on capital of existing data centers by optimizing around power constraints

Improving access to existing grid resources

[Google X's grid-optimization project, Tapestry](#), represents another pathway to solving the time-to-power bottleneck that can push data center developers toward large on-site BESS. Tapestry's GridAware and Grid Planning systems have several goals:

- Accelerate distribution inspections
- Improve transmission modeling confidence
- Give utilities and regulators earlier certainty that the grid can host new AI loads

Adopting these tools, or other means to reduce interconnection bottlenecks, erodes the value of using large BESS purely as a speed-to-power solution.

Utilities are using virtual power plants (VPPs) more often to provide peak-shaving, load-shifting, and local-capacity support by aggregating flexible behind-the-meter resources like batteries, HVAC systems, EV chargers, and heat pumps. These programs have grown to 19 gigawatts of capacity according to [Virtual Power](#)

[Plants: insights, Profiles and inventory](#) (PDF) from the Lawrence Berkeley National Laboratory. These programs give grid operators a way to get demand-side flexibility without relying only on large individual customers.

On-site BESS would continue to play a central role in ride-through, internal power-quality management, and backup. VPPs—many of which include distributed ESS—could become a parallel tool for providing certain grid-facing services and firm capacity.

As VPP participation grows, utilities may meet a greater share of their flexibility needs through aggregated distributed energy resources. This may make more power available to data center customers. On-site BESS would continue to play a central role in ride-through, internal power-quality management, and backup. VPPs—many of which include distributed ESS—could become a parallel tool for providing certain grid-facing services and firm capacity. The combination of data center demand-side and VPP supply-side solutions could upend our current notion of grid balancing.

Improvements to AI chip performance

AI chips turn on in synchronized bursts, so they need millisecond-level power smoothing. Thousands of GPUs ramp simultaneously, creating steep step-changes in power that UPS and grid interfaces must buffer. These problems stem from how manufacturers design GPUs and how GPUs operate, not from the processing workloads themselves. Operators won't compromise on the performance of processing workloads to solve power problems. But power availability could drive changes to chip design if the problem persists.

Data centers may require less BESS or UPS resources for stability if manufacturers improve chips to:

- Deliver more processing power per watt
- Need less simultaneous ramping
- Allow more modular training

Changes like these could reserve storage primarily for resilience or economic arbitrage instead of making up for instability in chip power. Load-smoothing BESS will see a strong window of opportunity for the next few years, but their long-term role is less certain as chip designs and workload orchestration improve. It's possible to imagine a world where tighter integration of energy storage into GPU racks unlocks greater performance. But it's also possible to imagine one where power smoothing happens elsewhere in the stack. The industry is evolving quickly and improving in many directions at once.

How can utilities and data centers form strategic partnerships?

In our exchange [Smart loads and smart strategies: Utility programs at the data center frontier](#), we explore

how utilities can engage local data centers as partners for energy efficiency, demand response, and load flexibility opportunities and programs.

Data center specific tariffs

Data center demand is driving the fastest electrical load growth US utilities have seen in decades. This poses a risk to utilities because accurately forecasting demand is how they plan investments and set rates. Utilities don't want to tell ratepayers they must pay for unusable infrastructure or that they won't get reliable electricity from the infrastructure they've already bought.

Many utilities and regulators are exploring higher fees, minimum load requirements, and contractual obligations to address the risks of new, large data center loads straining grid infrastructure. In July 2025, the Public Utilities Commission of Ohio [ordered AEP Ohio to create a new tariff](#) in its territory that's among the first of its kind in the US. The tariff will raise demand charges, increase collateral, and impose tighter contractual requirements to make sure that data center demand matches what operators promised. Dominion Energy [proposed a similar tariff structure](#) in September 2025. We expect more organizations to propose tariffs like these.

With tech cycles accelerating and tariff structures in flux, data center operators face rising planning risks. Rising risks will likely favor the investment case for BESS as a hedge against changes at utilities. Operators can view bring-your-own-capacity requirements, data center-specific tariffs, and on-site energy storage as coordinated system responses to large, inflexible loads rather than independent technology or policy strategies.

Future advances in chips, software orchestration, or grid coordination could reduce the need for batteries. But even if this occurs, AI data centers are unlikely to return to architectures that operate without fast, on-site energy storage.