



[www.pipelinepub.com](http://www.pipelinepub.com)

Volume 21, Issue 8

# How AI and Modular Design Are Powering the Next Generation of Data Centers

By: [Tod Higinbotham](#)

As data centers evolve into modular, AI-optimized hubs that improve scalability and efficiency, innovations like prefabricated modular designs, AI-driven automation, and advanced cooling systems are redefining the landscape.

As artificial intelligence (AI) workloads rapidly proliferate across the globe - creating unprecedented demand for power, cooling, and scalability - data center operators are reimagining their approach to data center design.



No longer confined to the one-size-fits-all architectures of the past, today's data centers are evolving to become AI-optimized, modular, purpose-built ecosystems that prioritize performance, flexibility, scalability, and energy efficiency.

This shift is occurring in response to the unprecedented demands that AI workloads impose on power, cooling, and scalability. As organizations race to deploy high-performance compute environments optimized for training and inference workloads, they find that traditional data center configurations are ill-equipped.

## The AI Load Challenge

The AI revolution introduced three distinct data center architectures. First, existing co-location facilities designed initially for general-purpose compute are now being retrofitted to accommodate GPU-intensive workloads, which require higher power densities and advanced cooling solutions. Second, new data centers are being designed from the ground up - some with more traditional compute workloads and many with AI-specific architecture in mind. The third is hyper-scale architectures.

[CoreSite's Chicago data center](#) (CH1) is a notable example of retrofitting. Originally designed as a telecommunications hub, CH1 has been upgraded to support modern high-density workloads. CoreSite incorporated redundant switchgear, spare generator positions, and enhanced cooling systems, enabling the facility to handle rack densities upwards of 40 kilowatts per cabinet with air cooling.

Colovore, a high-density hosting specialist, is [expanding its Santa Clara facility](#) to create a 9-megawatt data center optimized for AI workloads. This new facility will support rack densities of 50kW using water-chilled rear-door cooling units and can scale up to 250kW per rack with direct liquid cooling.

And, of course, there are the hyperscalers. Fueled by AI demand, hyperscalers, including Meta, Google, and Amazon, are [racing to build](#) some of the world's most significant data center clusters, which carry price tags of tens of billions of dollars per cluster.

## The Power Challenge

The differences between retrofitting and new builds are stark. In retrofitted environments, infrastructure must be adapted to meet higher power densities and address the dynamic nature of AI workloads. GPU clusters present a unique challenge in producing sharp, unpredictable power spikes. Power infrastructure and cooling systems originally optimized for flat CPU loads are now contending with a more volatile profile. In purpose-built facilities, designers can anticipate these challenges and engineer solutions that deliver optimized performance from the outset.

## Modularity: The Foundation of Flexibility

The shift to modular data center design marks a key innovation in the continuing transformation of data centers. While "modular" is used in different contexts, within infrastructure, it refers to the practice of preassembling sub-components of a data center - such as power distribution systems, cooling loops, and battery enclosures - off-site in controlled environments. These modules are then shipped to the deployment site, where they are integrated with minimal on-site labor.

Modular design accelerates deployment, improves build quality through factory-based testing, and lowers costs by reducing construction complexity. In practical terms, modular configurations enable phased buildouts, allowing operators to scale infrastructure incrementally in response to compute demand. This is especially critical as the AI hardware roadmap evolves at breakneck speed.

New modular shapes and configurations are being introduced to better align with AI-era infrastructure. For instance, instead of relying on legacy container-style layouts, new cabinet shapes are optimized for improved airflow, thermal efficiency, and space utilization within modular shells. Design teams are rethinking everything from battery enclosures to inverter cabinets to reduce floor space while boosting performance.

According to the [2025 Data Center Energy Storage Industry Insights Report](#) by Endeavor Business Intelligence, commissioned by ZincFive, 68% of 132 global industry professionals said they use modular power solutions. They plan to continue to do so, with 22% deploying them at all locations.

## Cooling Innovations for Thermal Realities

Liquid cooling is another innovation that has moved from the fringes of data center design into the mainstream. While reliable and well-understood, air-based cooling systems are becoming inadequate for the heat profiles generated by dense clusters of GPUs. Liquid cooling systems—whether direct-to-chip or immersive—can dramatically shrink the cooling infrastructure footprint while reducing energy consumption.

Liquid cooling systems also introduce a new layer of power dependency. Since they must operate continuously to prevent critical components from overheating, many operators are now provisioning uninterruptible power supplies (UPS) not only for compute workloads but also to ensure continuous cooling during grid outages. This dual-UPS architecture adds resilience and highlights how power infrastructure must evolve alongside thermal systems.

## The Rise of Pulse Loads and Dynamic Power Demand

Perhaps the most critical shift driven by AI workloads is the emergence of dynamic, pulse-based power profiles mentioned earlier. Unlike traditional compute, which exhibits stable and predictable loads, GPU-based training tasks often create synchronized power spikes as clusters of accelerators operate in unison.

This behavior creates new challenges for data center power design. Provisioning infrastructure to handle peak loads results in significant inefficiency, with excess capacity sitting idle during off-peak moments. To mitigate this, innovative energy storage technologies are being explored to absorb short-duration spikes and maintain consistent draw from the grid.

Recent engineering efforts have demonstrated how advanced battery chemistries such as ZincFive's [nickel-zinc \(NiZn\)](#) battery solutions can respond within milliseconds to these fluctuations, enabling data centers to smooth power profiles and reduce their reliance on overbuilt infrastructure. Since our battery solutions are very good at handling pulsing or transient power loads, we are working to set up systems to address this issue.

This evolution in energy storage will allow operators to match provisioned capacity to actual workload behavior better, significantly improving capital and operational efficiency.

## **Prefabrication and Integration: A New Way to Build**

Prefabrication is increasingly viewed as a cornerstone of next-generation data center builds. Rather than assembling all components on-site, operators now expect sub-assemblies - such as fully integrated UPS systems - to arrive tested, certified, and ready to deploy.

This approach improves quality control, accelerates construction timelines, and reduces the need for specialized on-site labor. It also promotes innovation, as manufacturers can iterate on design and functionality at the component level, ultimately delivering more compact and efficient systems. Data center operators

benefit from this model by gaining the flexibility to deploy consistent infrastructure across multiple locations, while also adapting to the specific needs of each site - whether that be high power density, geographic limitations, or sustainability targets.

## **Sustainability and Scalability**

Sustainability remains a defining consideration in modern data center planning. According to the same [industry insights report](#), nearly 90% of industry professionals cited sustainability as a key priority in 2025, up from 81% the previous year. Notably, the majority (72%) reported measurable cost savings from their sustainability efforts.

As operators seek to balance environmental responsibility with performance and profitability, modular configurations present a compelling solution. Factory-built modules reduce waste and emissions associated with on-site construction, and high-efficiency energy storage systems reduce reliance on diesel generators for backup power. The ability to scale incrementally also minimizes the risk of overbuilding and underutilizing infrastructure. When paired with sustainable nickel-zinc chemistry, these systems deliver environmental responsibility across the entire lifecycle, providing clean, reliable energy storage and immediate power from cradle to grave.

## **New Roles for Energy Storage**

One of the most notable shifts in energy strategy is the increasing importance of backup systems that do more than simply keep systems online during outages. Batteries are now being called on to perform new functions - including supporting cooling systems during outages, managing power transients caused by AI clusters, and even participating in grid-level stabilization programs.

Modular battery systems designed with high power density and fast discharge capabilities are particularly well-suited for this new role. These systems can be delivered fully prepackaged and tested, minimizing on-site integration time and aligning with the modular ethos of modern data center design.

The use of next-generation battery chemistries that require less cooling and fire suppression infrastructure, such as nickel-zinc, also contributes to lower total cost of ownership (TCO). By reducing the size and complexity of energy storage components, operators can improve rack density and reclaim valuable floor space.

## Looking Ahead: A Landscape in Flux

The pace of change in data center design continues to accelerate. Over the next three to five years, we can expect to see a proliferation of design approaches as operators balance innovation with the need for standardization and reliability.

While the journey may feel like the Wild West today, the destination is clear: data centers that are faster to build, easier to maintain, and more responsive to the needs of high-performance workloads.

Meeting this future will require deep collaboration across the supply chain. Energy storage, thermal systems, compute hardware, and software must be harmonized to deliver the performance, reliability, and sustainability that the next generation of applications demands. Modular design isn't just a trend—it's the blueprint for data center innovation in the AI era.

Not for distribution or reproduction.