WHITEPAPER | AI DRUPS

Mission Critical Facilities
Artificial Intelligence and Dynamic Rotary UPS Systems

Introduction

For more than 60 years, Dynamic Rotary Uninterruptible Power Supply systems (DRUPS) have played a key role in protecting critical processes. As data centers expanded particularly during the dot-com boom of the mid-1990s, it was only natural that wlarge-capacity Dynamic Rotary UPS systems were adopted to support their growing power demands. By the early 2000s, the surge in internet adoption, cloud computing, and global business operations fueled an unprecedented demand for scalable and reliable IT infrastructure. This led to the rise of purpose-built data centers: massive facilities designed from the ground up to host IT infrastructure at scale. Over time, these evolved into today's mega and hyperscale data centers,which require even greater UPS power capacity.

Dynamic Rotary UPS systems have been deployed in hundreds of data centers worldwide and continue to operate reliably, often for more than 25 years without replacement. Their nominal unit power has steadily grown to 3600 kVA, making them suitable not only for hyperscale data centers but also for other critical applications. A key advantage of Dynamic Rotary UPS technology is its ability to support any load, independent of its profile. This flexibility has established Dynamic Rotary UPS systems as a trusted solution for protecting vital processes across industries such as semiconductor manufacturing, telecommunications, and critical transport infrastructure.

HITEC Diesel Rotary UPS (DRUPS)

A traditional static UPS system with batteries typically consists of a rectifier, batteries, and an inverter. It provides temporary backup power through its batteries before transferring the critical load to a standby generatorduring extended outages.

In contrast, the HITEC Dynamic Rotary UPS (DRUPS) integrates all these functions into a single, compact and efficient system comprising a synchronous generator, kinetic energy module, and a diesel engine. When utility power is interrupted, the system draws stored kinetic energy from its Kinetic Energy Module (KEM) to maintain power to the critical load via the synchronous generator until the integrated diesel engine takes over.

Continuous Cooling

As rack power density in data centers increases, the demand for continuous and reliable cooling of IT servers becomes critical. To address this, many operators have connected their entire cooling systems to the no-break output of the Dynamic Rotary UPS units.
This approach allows the entire cooling systems, including motor and compressor loads, to receive an uninterrupted power supply and thereby, eliminating concerns about thermal storage, complex control schemes, or thermal runaway during utility outages. The Uptime Institute has also recognized and recommended this method in its Accredited Tier Designer Technical Paper Series: Continuous Cooling.

Operators can use Dynamic Rotary UPS to support any type of load. This allows them to replace multiple single-purpose systems like static UPS units, batteries, generators, and transfer switches with one integrated solution. It simplifies the electrical infrastructure for large-scale systems. This approach also streamlines the cooling setup, ensuring uninterrupted cooling during power outages. This advantage is especially valuable today as data center construction faces new challenges. AI is changing how heat is managed in these facilities, and many operators are rethinking their power strategies because their current architecture can't keep up.

The traditional data centers server racks consume an average of 7 kW. But when integrating AI hardware, consumption can surge to 50–100 kW or more. This four- to fifteen-fold increase in power consumption generates significant additional heat, making efficient cooling critical for both uptime and performance. To meet these demands, data centers must shift from brute-force heat removal methods, which cool the entire rack, towards targeted solutions that address the primary heat source: the chips themselves. Direct-to-chip OTC liquid cooling has emerged as the most effective solution. By cooling chips directly, data centers can operate more efficiently, optimizing both energy and water consumption. This approach can reduce cooling-related energy consumption by nearly 20%, freeing more resources for computing needs, while also allowing higher server density per rack.

As a rule of thumb only around 60-70% of the heat rejected by a GPU AI chip can be captured by DTC cooling, the remainder is rejected into the server room; and we know that air cooling is sufficient for Server Racks up to 30kW. So, for dense AI applications, both effective air (room) and liquid (chip) cooling is essential. Add to this the aggressive nature of an AI Learning load profile, a cooling challenge is created that must ensure the sufficient delivery of cooling during load peaks and sufficiently limit cooling during load sags.
To prevent the potential for over or under cooling of the GPU's and the environment is then further complicated by the necessity to mitigate the effects of any power outages. Inevitably, this leads to a highly complex cooling system comprising many separate components, numerous discrete controls needing multiple changes in operational state, and the development of an overarching control system to bind everything together to guarantee safe operating temperatures of around 40degC are not exceeded.

One option to address this thermal intertia shortfall is to install massive chilled-water tanks for AI clusters, along with additional heat exchangers, chillers, pumps, and redundant cooling infrastructure. While this approach can maintain water flow during restart periods, it requires significant capital investment, space, energy, and maintenance. Comprising numerous individual components and of course, such multi layered mechanical solutions require significant complex engineering to make the system ready for different failure scenarios. Planning for worst-case thermal scenarios is a crucial part of this design.

Direct liquid cooling

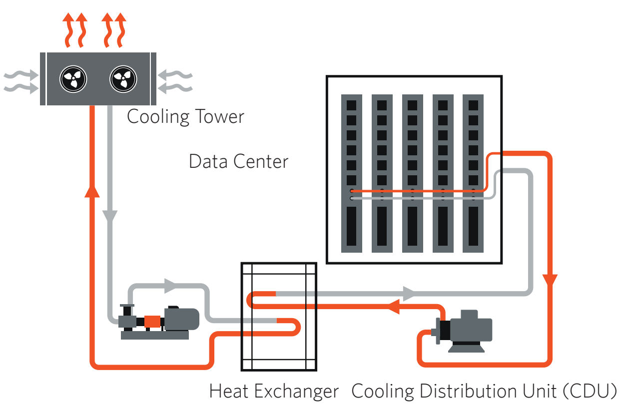

A simpler and more reliable alternative is to connect the cooling system directly to the no-break output of a Dynamic Rotary UPS. Because Dynamic Rotary UPS systems are unaffected by load profiles, both IT and cooling loads can be protected by a single system. This ensures that in the event of a power outage, IT equipment and its cooling remain uninterrupted, avoiding the dangerous scenario of servers continuing to run while cooling stops. For AI environments, Dynamic Rotary UPS technology therefore offers a far more straightforward resilient power strategy resulting in higher uptime & improved reliability.

Understanding AI Power Requirements

One of the greatest challenges in the transition to AI and high-performance computing (HPC) is adapting power supply systems to meet rapidly increasing demands. Traditional power distribution systems were never designed to support the extreme requirements of AI-driven workloads. While current AI racks typically consume 50–100 kW, data center operators are already planning for densities as high as 150 kW per rack, with projections of up to 500 kW on the horizon. This represents a dramatic shift in data center design philosophy.

As a result, attention is now focused on building robust power infrastructures capable of supporting this new paradigm. As a result, circuit breakers, busbars, and PDUs are sized with little or no headroom in AI facilities. Engineers must assume close to 100% load scenario. This affects everything from transformer capacity (more MWs needed) to cooling plant steady-state load (running near maximum). It also means that the entire electrical infrastructure operates under heavier continuous load than in a typical data center. The power infrastructure must be robust to handle this strain continuously. These AI loads put the critical power equipment, like the UPS, under unprecedented stress. Dynamic Rotary UPS units provide a compelling solution with nominal unit power ratings up to 3600 kVA, making them ideal for large-scale (high-power) centralized deployments. In fact, AI data centers feature electrical systems akin to those of industrial processes and power plants, not just IT server rooms. And this is the environment, where Dynamic Rotary UPS systems have been dominating already for decades, securing high-value industrial installations of tens of MWs, often centralized as, for example, medium voltage UPS systems. Consequently, unlike decentralized systems composed of numerous smaller UPS units and complex switching equipment, centralized Dynamic Rotary UPS systems simplify these high-power electrical infrastructures while enhancing reliability.

AI workloads also introduce highly dynamic and unpredictable power fluctuations, with rapid surges and sags in power demand. Traditional static UPS systems, designed for more stable and predictable limited loads, struggle under these conditions.

Fluctuating power demands accelerate thermal cycling, shorten the lifespan of electrical components, and increase the risk of service interruptions. Many AI data centers therefore prioritize critical systems for UPS and allow non-critical compute loads to ride on utility power with generator backup only leaving the sensitive AI hardware unprotected against voltage sags or transients from the grid.

To maintain power quality, stability, and reliability, AI data centers therefore require large power “AI-tolerant” UPS systems engineered to handle such load swing variability. Dynamic Rotary UPS technology, combining the protection of a flywheel UPS with integrated diesel generator back up, is such a large power AI-tolerant UPS system, and is already proven in managing fluctuating mechanical loads in industry. Its integrated oversized alternator can absorb large load steps, inrush currents, and even short-circuit currents up to 10–20 times nominal values. Moreover, its ability to support any load profile makes it uniquely suited to AI’s unpredictable power patterns.

This is not just theoretical: for over two decades, Dynamic Rotary UPS systems have been deployed not only to support high-power industrial critical processes, but also to support critical oscillating loads such as pulse radars a bad profile that closely resembles the synchronized bursts large of AI language model workloads. This long-standing operational track record demonstrates that Dynamic Rotary UPS systems are well-positioned as a proven technology to support the new demands of AI computing.

Selecting an Uninterruptible Power Supply

Selecting the right uninterruptible power supply (UPS) is critical to safeguarding a data center's infrastructure and ensuring continuous operations. Traditionally, designers and operators have assumed that UPS support applies exclusively to IT loads. However, Dynamic Rotary UPS systems offer a broader and more flexible approach.

With unit capacities reaching up to 3600 kVA, Dynamic Rotary UPS systems not only simplify high-power data center infrastructure but also support additional systems beyond IT equipment. In practice, this means that both highly fluctuating AI and High Performance Computing workloads, as well as entire cooling systems, including mechanical loads, can be continually powered and protected by a single Dynamic Rotary UPS.

This integrated strategy eliminates the risks associated with traditional separated designs, where IT loads remain on-line during an outage, but the cooling system do not. By consolidating IT and cooling onto the same no-break output, Dynamic Rotary UPS ensures that both remain operational under all circumstances. This creates a unified power train capable of handling the full spectrum of demand, but also avoids rapid transitions from near-zero to full load within milliseconds, something static single purpose UPS systems have to face and struggle to achieve.

Such an approach represents a significant step forward in power strategy for AI and high- performance computing environments. It not only enhances reliability but also simplifies infrastructure design, reducing points of failure while improving overall efficiency. Several data centers have already implemented this solution successfully, providing real-world validation for its effectiveness in next-generation facility design.

Additional Advantages of Dynamic Rotary UPS

Beyond their ability to protect fluctuating IT and cooling loads, Dynamic Rotary UPS systems offer several other advantages over traditional static UPS designs that rely on batteries and gensets.

One key difference is the use of an integrated kinetic energy storage module instead of batteries. As an electrical machine, the kinetic module removes the risk of fire associated with lithium-ion, reduces cooling requirements (lower losses), and heavy maintenance associated with batteries. Currently the Kenetic Energy Storage within the Dynamic Rotery UPS provides unrivalled power density, recharges within minutes, has no limitations on the number of charge–discharge cycles, and operates efficiently for at least 25 years.

The integrated engine, equipped with a dedicated redundant starting system, delivers the most reliable engine start capability in the industry. Combined with single-shaft technology, this design ensures exceptional autonomy. Dynamic Rotary UPS systems also support fuel- and emission-reduction measures, including the use of HVO (Hydrotreated Vegetable Oil), start delay optimizations, load-dependent redundancy, and “fuel saving” modes.

From an environmental perspective, Dynamic Rotary UPS systems are highly sustainable. All materials are fully recyclable, no toxic or chemical components are used, and the carbon footprint of the kinetic module is several times lower than that of batteries found in traditional systems.

Finally, Dynamic Rotary UPS systems deliver significant efficiency and cost advantages. By reducing the Comparative UPS footprint by 40–50%, they optimize space utilization while lowering both capital and operational costs, resulting in a far lower total cost of ownership (TCO) for high-power data center infrastructure.

Conclusion

The rapid adoption of AI and high-performance computing is fundamentally reshaping data center design. Power densities that once averaged 7 kW per rack are now climbing to 100 kW and beyond, creating unprecedented challenges for both power and cooling infrastructures. Traditional static UPS systems and battery-based designs, while effective in legacy environments, need adaptation to handle the scale, variability, and resilience required by modern AI-driven workloads.

Dynamic Rotary UPS systems offer a proven and forward-looking solution. Their ability to support any load profile, whether stable, highly dynamic, or mechanical, makes them uniquely suited to protect both IT and cooling systems under a single, centralized continuous power delivery architecture. By eliminating batteries, integrating kinetic energy storage, and incorporating highly reliable engine technology, Dynamic Rotary UPS systems deliver unmatched resilience, sustainability, and efficiency.

For operators, the benefits are clear: simplified infrastructure, reduced footprint, lower total cost of ownership, and improved environmental performance. For mission-critical industries, from hyperscale data centers to semiconductor manufacturing and critical infrastructure, Dynamic Rotary UPS provides the reliability and adaptability needed to meet tomorrow's challenges.

As AI continues to drive exponential increases in power demand, the need for robust, scalable, and efficient power protection will only grow. Dynamic Rotary UPS systems are not just an alternative to static UPS, they represent the next generation of power infrastructure, enabling data centers to thrive in an era defined by digital transformation and artificial intelligence.

HITEC Power Protection BV
Bedrijvenpark Twente 40
7602 KB Almelo
The Netherlands
Phone +31 546 589 589

hitec-ups.com

**CONTINUOUS POWER
IN YOUR CONTROL**