Stabilizing AI Power at Scale with e-POD: A Layered Critical Power Architecture

AUTHOR
Pierre-Adrien Bel
Product Manager - Large Diesel Generators

INTRODUCTION

Artificial intelligence workloads are reshaping the electrical behavior of data centers. Unlike conventional IT loads, AI training and inference clusters are driven by synchronized graphics processing unit (GPU) operation that produces rapid, repetitive, and high-amplitude variations in power demand. These fluctuations occur on timescales of tens of milliseconds—well below the response capabilities of traditional mechanical power systems.

Research published by Microsoft and NVIDIA shows that large-scale AI training workloads operate in synchronized cycles of compute-intensive and communication-intensive phases, causing rapid power ramp rates and repetitive load oscillations (see *Figure 1*).

When thousands of GPUs transition between these phases in lockstep, the resulting coherent power fluctuations can occur at frequencies that are poorly damped by the electrical grid and may propagate through the power train if not actively managed.

This paper describes system-level architectural principles. Detailed control logic, sizing methodologies, and operational strategies are intentionally outside its scope.

rehlko
AI Ready

Figure 1

Microsoft illustration of power readings from an at-scale training job on DGX-H100 racks

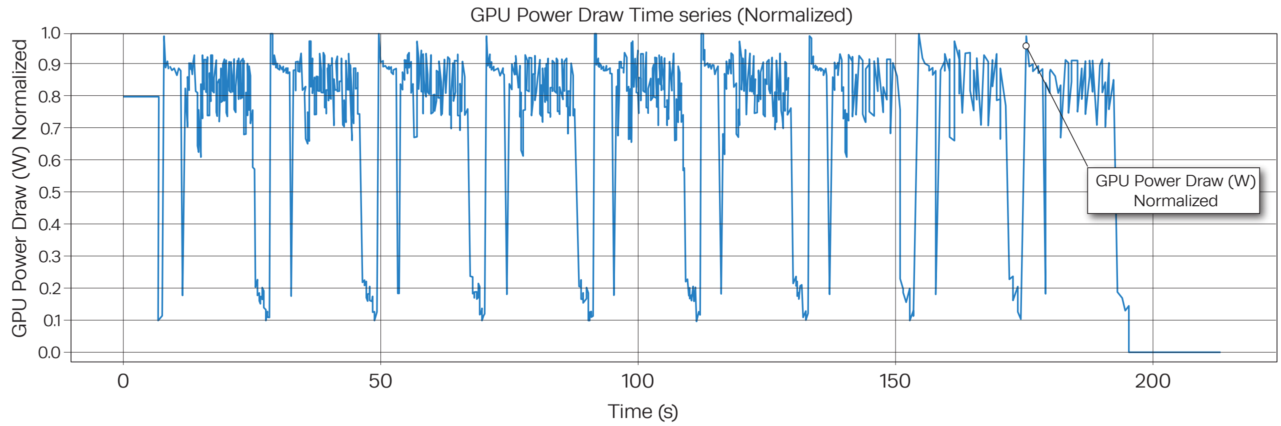

Figure 1 – Simplified data center power train

Illustration of the electrical power train showing the flow of energy from the utility grid or on-site generation, through medium- and low-voltage switchgear, transformers, and power electronic interfaces like uninterruptible power supplies (UPS), to the IT load.

THE ELECTRICAL POWER TRAIN

These characteristics challenge long-standing assumptions about power stability in data center environments. Architectures designed primarily for steady or slowly varying loads may be insufficient when confronted with the highly dynamic behavior of AI workloads. As a result, AI-driven data centers require additional system-level mitigation strategies capable of addressing power volatility across multiple timescales, from milliseconds to days.

This shift necessitates a holistic approach—one that treats power infrastructure not as a collection of isolated components, but as an integrated electrical system designed around the physical characteristics of modern AI computation.

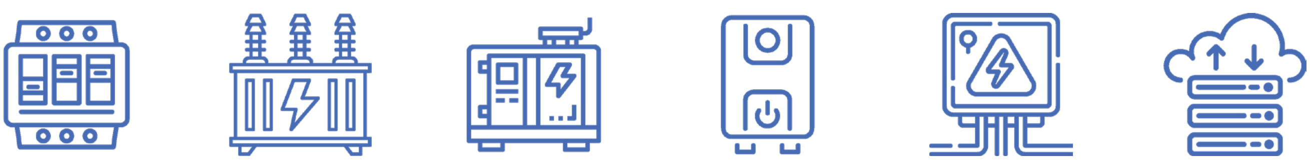

Medium- and low-voltage switchgear (MLVS)

Transformers (MV/LV, isolation, step-down)

On-site source for production or backup power

Power-electronic interfaces (UPS, converters, inverters)

Distribution paths (busways, PDUs, remote power panels)

End loads (GPU racks, servers, networking equipment)

A MULTI-LAYERED POWER ARCHITECTURE

Joint research by Microsoft, NVIDIA, and OpenAI confirms that AI workloads introduce high-frequency, repetitive electrical oscillations that cannot be addressed by mechanical generation systems alone. The inertia and electromagnetic response limits inherent to rotating machines constrain how quickly they can react to such dynamics. As a result, electrical instability originating at the load level can propagate upstream if not actively mitigated.

Stabilizing AI power therefore requires multiple layers of response, each aligned with a distinct timescale of system behavior:

- Sub-second electrical buffering to absorb fast load oscillations.
- Continuous power conditioning and electrical isolation to maintain voltage and frequency stability.
- Long-duration energy sources to ensure resilience during extended grid interruptions.

e-POD ARCHITECTURE

Rehlko's e-POD architecture is structured explicitly around this multi-timescale challenge. Rather than relying on generator oversizing or localized mitigation techniques, e-POD integrates three complementary functional layers, each optimized for a specific role within the overall power system:

- **StoreModule:** battery-based energy buffering for millisecond-scale stabilization. It is the only element that contains dedicated energy storage and is designed to provide fast-acting electrical buffering for AI-induced load dynamics.
- **ShieldModule:** active power conditioning and electrical isolation through advanced UPS technology. It contains no long-duration energy storage. Its role is strictly limited to power electronic conversion, regulation, and electrical isolation between upstream sources and downstream loads.
- **GenModule:** long-duration generation for sustained resilience during outages.

Together, these layers form a coordinated electrical architecture capable of supporting highly dynamic AI workloads while limiting the propagation of destructive electrical variability into upstream infrastructure.

LAYER 1 – STOREMODULE

STABILISING AI VOLATILITY AT THE FASTEST TIME SCALE

Current research indicates that battery-based energy storage, when combined with power electronics, is among the most effective approaches for managing millisecond-scale load fluctuations introduced by AI workloads. (see *Figure 2*)

Figure 2

NVIDIA illustration of synchronized GPU compute/communication cycles causing rapid, repetitive power oscillations in AI training workloads.

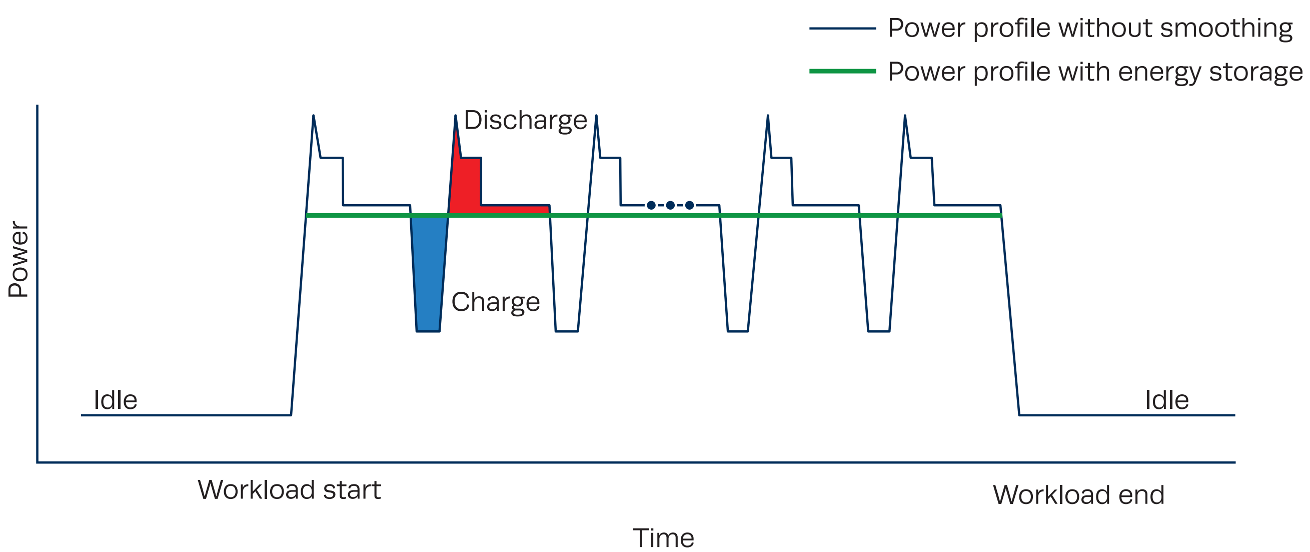

The StoreModule is not intended to act as a long-duration energy source. Its function is to absorb and release power at very high speed in response to load ramp rates (dP/dt), reshaping fast electrical transients into a stable upstream demand profile rather than supplying sustained energy.

Within the e-POD architecture, the StoreModule provides fast-acting energy buffering in support of system-level power stabilization. Rather than acting solely as a backup energy resource, it operates in coordination with the ShieldModule to absorb and release energy in real-time, mitigating rapid and repetitive load variations at their point of origin.

By buffering short-duration power fluctuations and reshaping them into a more consistent demand profile, the StoreModule limits the propagation of high-frequency electrical dynamics beyond the load environment. This behavior reduces stress on both downstream power conditioning equipment and upstream infrastructure, contributing to more predictable and stable overall system operation.

LAYER 2 – SHIELDMODULE

CONDITIONING AND ISOLATION THROUGH POWER ELECTRONICS

Located at the interface between a highly volatile load environment and long-duration energy sources, the ShieldModule serves as the primary stabilization layer within the e-POD architecture. Built around Rehlko's MF Series UPS, it continuously maintains voltage and frequency stability under rapidly changing operating conditions.

Within the e-POD architecture, the ShieldModule provides no dedicated energy storage. Its function is limited to power-electronic conversion, conditioning, and electrical isolation. Fast energy buffering is intentionally handled by the StoreModule, while long-duration backup energy is supplied by the GenModule, allowing each layer to be aligned with the timescale it is best suited to manage.

Rather than operating as a passive protection element, the ShieldModule actively regulates power flow between the electrical grid—or the GenModule during grid interruptions—and the StoreModule. In doing so, it electrically decouples upstream infrastructure from the fast, repetitive disturbances produced by AI workloads, while delivering a stable and predictable power profile to the IT load.

In this role, the UPS extends beyond traditional backup functionality and operates as a real-time power control interface, governing how energy is conditioned, isolated, and delivered across the system.

Traditional UPS battery systems are designed to provide short-duration ride-through during grid interruptions and generator transitions.

They are not optimized for repeated, millisecond-scale charge and discharge cycles associated with AI load oscillations.

Within the e-POD architecture, the ShieldModule therefore does not rely on its own energy storage for load smoothing. Fast and repetitive AI-driven power fluctuations are intentionally handled by the StoreModule, avoiding excessive cycling and premature aging of UPS batteries.

StoreModule and AI Loads: Clarifying the Role

Recent research conducted by Microsoft and NVIDIA highlights the importance of fast-acting energy buffering to manage the highly dynamic electrical behavior of large AI training clusters. While rack-level energy storage and local buffering solutions are effective at absorbing very short-duration transients, they do not fully address the cumulative, system-level effects generated by massively synchronized GPU loads.

Within the e-POD architecture, centralized energy storage is therefore not positioned as a simple backup capacity, but as a system-level stabilization layer. By electrically decoupling volatile AI loads from upstream infrastructure, this storage converts rapid, repetitive power variations into a stable and predictable demand profile, reducing integration risk and protecting both grid interconnections and generator sets.

This approach complements rack-level stabilization strategies and enables AI-oriented data centers to scale while maintaining electrical stability across all relevant timescales.

LAYER 3 – GENMODULE

LONG-DURATION RESILIENCE WITHOUT MECHANICAL STRESS

Generators remain a critical component for supplying long-duration power during extended grid interruptions. However, the rapid and repetitive load variations associated with AI workloads occur on timescales well below the inherent response limits of mechanical generation systems.

The operational impact of directly exposing generators to such high-frequency load profiles is not yet fully characterized. This operating regime is relatively new and differs materially from traditional data center load behavior, resulting in limited industry experience and long-term data validation.

By addressing fast electrical dynamics upstream through dedicated buffering and conditioning layers, the e-POD architecture prevents AI-induced load volatility from directly interacting with generator control systems. This allows generators to operate within well-understood, validated performance envelopes without reliance on oversizing or mechanical inertia as a mitigation strategy. See *Figure 3*.

Figure 3

Response Time Comparison Across Power Systems	
SYSTEM	**TYPICAL RESPONSE TIME**
ICE Generator	200 – 450 ms
AI Load	10 – 200 ms
MF Series UPS	2 – 4 ms

WHY e-POD IS THE CORRECT ARCHITECTURE

The electrical challenges introduced by AI workloads are no longer theoretical. Extensive analysis from Microsoft, NVIDIA, and joint Microsoft–NVIDIA–OpenAI research documents a convergence of fast load transients, high-amplitude power swings, and frequency-sensitive oscillations that originate directly from synchronized GPU operation.

Figure 4 illustrates how these observed behaviors map to the layered structure of the e-POD architecture. Millisecond-scale fluctuations resulting from GPU lockstep activity cannot be addressed by mechanical systems alone and must be mitigated electrically at the fastest timescale. Within the e-POD, this function is addressed by the StoreModule, operating in coordination with the ShieldModule, before such dynamics can propagate upstream.

Larger power variations associated with alternating compute and communication phases further emphasize the need for close coordination between energy buffering and active power conditioning. Here, the combined action of the StoreModule and ShieldModule reshapes these fluctuations into a more stable and predictable demand profile.

Frequency-sensitive oscillations identified in both Microsoft and arXiv studies require electrical isolation rather than additional energy buffering. This role is fulfilled by the ShieldModule, which maintains voltage and frequency stability independently of upstream grid or generation behavior.

Long-duration outages remain the responsibility of conventional power generation. Within the e-POD architecture, the GenModule provides extended resilience while operating behind electrically stabilizing layers that limit exposure to AI-induced power volatility.

Taken together, *Figure 4* illustrates that each identified AI-driven electrical challenge corresponds to a specific architectural function. The e-POD is therefore not a collection of discrete components, but a system-level architecture intentionally designed to address the full range of electrical behaviors introduced by AI workloads.

Figure 4

POWER CHALLENGE	AI SOURCE DESCRIPTION	e-POD SOLUTION
Millisecond fluctuations	GPU lockstep oscillations (NVIDIA)	StoreModule + ShieldModule
High-amplitude swings	Compute/communication phases (Microsoft	StoreModule + ShieldModule
Frequency-sensitive oscillations	Harmonic bands risk (Microsoft & arXiv)	ShieldModule
Long-duration outages		GenModule

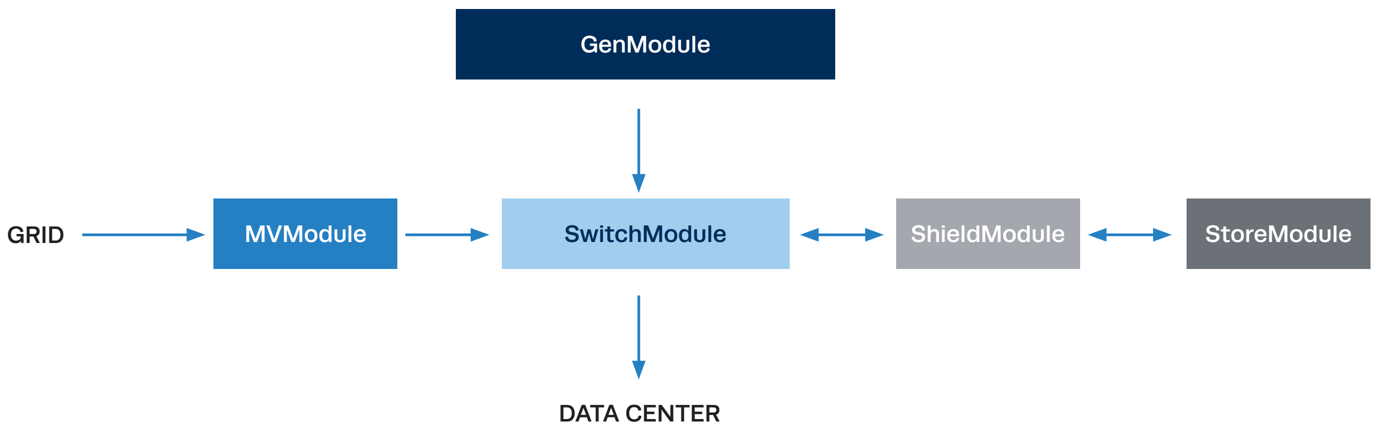

CONCLUSION

AI workloads fundamentally redefine the concept of power stability in modern data centers. The synchronized, high-frequency behavior of GPU clusters introduces electrical dynamics that exceed the physical response limits of traditional power systems when considered in isolation. Assumptions that once held true—namely steady or slowly varying IT loads—are no longer valid in AI-driven environments.

Research conducted by Microsoft, NVIDIA, and the joint Microsoft–NVIDIA–OpenAI collaboration demonstrates that AI-related power challenges span multiple timescales simultaneously. Millisecond-level oscillations, high-amplitude power swings, frequency-sensitive interactions, and long-duration outage requirements must be addressed in a coordinated manner. Approaches based on isolated components or incremental oversizing fail to capture these interactions and may introduce additional operational risk.

The e-POD architecture addresses this reality through a system-level design approach. By combining fast-acting energy buffering, active power conditioning and isolation, and long-duration generation within a coordinated electrical platform, each technology is aligned with the timescale it is physically capable of managing. Electrical volatility is mitigated at the point of origin, power quality is continuously enforced, and generators are reserved for sustained energy delivery.

Critically, this approach treats critical power infrastructure as an integrated system rather than a collection of individual products. Through architectural clarity and factory-level system validation, e-POD reduces integration risk and provides predictable, repeatable system behavior in deployment.

As AI workloads continue to scale, data center power architectures must evolve accordingly. Meeting this challenge requires a shift away from component-centric designs toward solutions engineered around the dynamic electrical behavior of AI computation.

The e-POD architecture represents such an evolution—designed to maintain electrical stability across timescales ranging from milliseconds to days and to support the continued expansion of AI-driven infrastructure. *https://www.wbpsltd.co.uk/projects/basecore-pods/*

Observed AI-related electrical behaviors and architectural responses described in this document are based on currently available research and industry observations. As AI workloads and operating practices continue to evolve, further evaluation and validation may be required to fully characterize long-term system impacts.

REFERENCES

- *[2508.14318] Power Stabilization for AI Training Datacenters*
- *Power stabilization for AI training datacenters | Microsoft Community Hub*
- *Powering Up AI: Microsoft, Nvidia, and OpenAI Tackle Datacenter Power Swings | AI News*
- *How New GB300 NVL72 Features Provide Steady Power for AI | NVIDIA Technical Blog*

ABOUT THE AUTHOR

Pierre-Adrien Bel is a Product Manager responsible for large diesel generator sets above 700 kW. Holding a degree in mechanical engineering, he brings extensive experience in the power generation industry.

He began his career at Rehlko in engineering in 2008, overseeing generator set installations, before moving into project management for complex, custom projects. He then worked as a Sales Engineer for special applications. Since 2021, he has been part of the Product Management team.

He is also the author of Product Environmental Profiles (PEP).

His areas of expertise include codes and standards, generator set packaging and installation, emissions optimization, sustainable solution development, and decarbonization strategy.

ABOUT REHLKO

A global leader in energy resilience, Rehlko delivers innovative energy solutions that sustain and improve life across home energy, industrial energy systems, and powertrain technologies with control, resilience, and innovation. Leveraging the strength of its portfolio of businesses — Power Systems, Clarke Energy, Home Energy, and Engines—and its more than a century of industry leadership, Rehlko provides power where and when the grid cannot. Rehlko goes beyond function and individual recovery to create better lives, communities, and a more durable and energy-resilient future.

rehlko