# Let’s talk behind-the-meter (island) power for AI data centers

## The plain and simple reality.

Author:  William Langley, Main Board Director, Langley Holdings plc

Running a data center off-grid (in true behind-the-meter island-mode), introduces a critical challenge that has caught many operators off guard: the volatility of AI transients.

Unlike conventional data center demand, which is typically steady and predictable, AI workloads behave very differently. They arrive in bursts. When a new round of training or inference jobs begins, large clusters of high-performance compute servers can draw a surge of power in just milliseconds.

When those jobs finish or are throttled back, that additional demand can disappear just as quickly.

In real AI data center operations, these fluctuations are not marginal. Load demand can step from 80% → 50% →90%, which in a 100MW+ class plant means tens of megawatts of load appearing or disappearing almost instantaneously.

On a grid-connected site, the public grid acts as a vast shock absorber — immediately delivering extra power when demand spikes and readily absorbing it when demand falls.

An islanded power plant has no such cushion. The on-site power plant must cope with these sudden load steps itself. The reality? Reciprocating engines (and even more so turbines) have physical “ramp rate” limits on how quickly they can ramp power up or down. At typical operating points (e.g. at 90% load) a modern medium-speed gas-fired engine can only tolerate an instantaneous step of about 1% of its rated load without frequency dipping below acceptable limits. Beyond that small step, further load must be taken up gradually, far slower than the rapid multi-MW surges created by AI workloads.

**Multiple gensets**

When the load demand rises or drops faster than the island plant can respond, frequency and voltage become unstable. Left unchecked, this turbulence drives inefficiencies, stresses equipment, and in extreme cases can lead to plant trips.

For independent power providers (IPPs) aiming to guarantee five-nines availability, that level of instability and unreliability simply isn’t acceptable.

## The “solutions” have been limited.

Facing these sharp, AI-driven load steps, operators have turned to two broad groups of measures — on the power-supply side and on the data-center side.  Each brings partial benefit, but none on its own can truly protect an islanded power plant.

## Power-side approaches.

### Oversize the power plant:

The traditional response is to install more engines or turbines so there is plenty of spare capacity, in the form of “spinning reserve”, ready to respond.

It works to some extent but comes at a high price: higher CapEx, more maintenance, more fuel burn, higher installed emissions — and much of that capacity sits idle most of the time, wasting headroom. It’s a blunt tool that does not address the real problem: the speed of the ramp.

## Deploy BESS:

Battery Energy Storage Systems (BESS) are often added at the plant level as a buffer. They can smooth modest spikes and dips by briefly supplying or absorbing power. But AI transients demand discharging and charging in milliseconds, and that's where batteries struggle with C rate limits and the need to oversize to enable absorption at a multi-MW scale. Their energy is stored in electrochemical reactions that struggle to move ions fast enough for these very high-frequency, high-magnitude fluctuations.

Forcing them to do so heats the cells and accelerates degradation, shortening their lifecycle. Even with large installed capacity, BESS is poor at supplying the short-circuit current needed for selective fault-clearing in an islanded grid.

So BESS can help moderate short-term fluctuations, but in an islanded power plant their chemical and endurance limitations and oversized space-hungry structure make it unsuitable as the core stabilization layer for fast AI workloads.

# Data-side approaches.

## Centralised UPS:

The usual response is to install larger, more resilient (and more costly) centralised UPS systems.
Major OEMs such as Schneider Electric, Vertiv, and Eaton offer "AI-ready" UPS platforms designed to tolerate fast, short-lived load ramps and improve resilience for sensitive IT racks. However, they stabilise the load, not the power source.

They cannot stop violent multi-megawatt transients from propagating upstream to the power plant, and their batteries degrade faster when forced to respond repeatedly to rapid steps.

## Raise the power floor:

By keeping the IT load artificially elevated, often through orchestration-level software scheduling, operators shrink the gap between idle and peak demand. This approach softens the deepest trough-to-peak ramps but comes with trade-offs:
– The baseline power plant load, and therefore fuel burn and OpEx, remain higher even during low-compute periods.
– Peaks can still rise above the raised floor, meaning reserve power must still be kept online.

## Rack-level smoothing:

The latest advances aim to smooth AI transients directly at the rack level, before they ever propagate upstream to the power source.

New platforms such as NVIDIA's GB300 NVL72 integrate capacitor-based energy storage, firmware-controlled ramp-rate management and a 'burn-mode' discharge function, potentially cutting the amplitude of instantaneous spikes by up to ~30% and softening short bursts that would otherwise appear as abrupt load steps to the supply.

This is a meaningful step forward, particularly for grid-connected data centres. However, even after this smoothing, each rack can still jump by tens of kilowatts — and when multiplied across thousands of racks, that residual variation translates to multi-megawatt load steps. A public grid can readily absorb those steps, but an islanded power plant cannot without dedicated stabilization at the source.

# There's a smarter way.
# SHIELDX™ by Piller — Dynamic Power Stabilization (DPS).

Deployed in mission-critical infrastructures worldwide (including the Federal Reserve, the New York & London Stock Exchanges, and the CIA), SHIELDX™ brings together two complementary battery-free technologies that work in tandem at the plant level to address the very behaviors that make AI-driven loads so disruptive.

## First — the high-inertia stabilizer (not UPS).

This is the backbone of the system. It's a continuous electromechanical device, always spinning in synchronism with the power source. It doesn't have to "wake up" or wait for software logic to act. That means it reacts in sub-milli-seconds, instantly resisting sudden load changes.

It acts like a shock absorber between the power plant and the data-center load, smoothing out micro-spikes that would otherwise appear as voltage flicker, keeping the voltage profile steady and holding system frequency typically within ±1% even during sharp ramp periods. Just as importantly, it delivers the high short-circuit current needed for selective fault-clearing.

## Second — the bi-directional power-exchange modules.

These handle the large, rapid ramps profiles that engines and turbines cannot follow in time.

Continuously online, not dependent on battery inverter switching, they respond in sub-millisecond timeframes, instantly injecting or absorbing power.

When AI workloads suddenly surge, or when a generator/turbine unexpectedly drops offline, the modules inject active power almost instantaneously, holding frequency and voltage steady and "buying time" for the remaining plant to ramp-up at its optimal ramp rate, and/or for standby units to join the bus.

The reverse is just as important.

When AI clusters throttle back and load collapses, these modules instantly absorbs the surplus energy that the generation fleet continues to produce during their ramp-down. Think of it as a vehicle suddenly downshifting from sixth gear to third: without this buffer, that excess energy would have nowhere to go, risking overspeed and frequency spikes.

Instead, it's kinetic energy storage modules momentarily absorbs the excess power and then releases it smoothly so that, from the data-centre side, the power supply remains continuous and stable.

## The outcome.

The power plant sees a smoother, gradual ramp profile, allowing the power generation fleet to operate near their optimal steady-state load instead of chasing spikes and troughs. This reduces wear, lowers unplanned maintenance, improving fuel economy, and maintaining system frequency within ±1% — shielding the fleet from impossible load profiles, instability, and stress, even during sudden large AI-driven load step

The data center continues to generate sharp AI load profiles but receives a clean, stable supply — shielded from frequency dips and voltage flicker — so that UPS and IT loads perform as if powered by a premium-quality power source.

SHIELDX™ protects both sides of the power chain, sustaining this level of performance for decades without degradation. It can be readily deployed with any power source — typically high-speed or medium-speed reciprocating engines, or turbines.

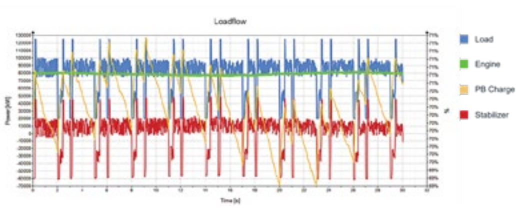


**ShieldX in action**

## Langley Power Solutions.

The Langley Power Solutions Division brings together three globally respected OEMs: Bergen Engines, Marelli Motori, and Piller Power Systems.

When Piller's SHIELDX™ is paired with Bergen Engines' lean-burn, medium-speed gas-fired engines and Marelli high-efficiency alternators, the result is a complete power plant designed to deliver stable, efficient and fast-to-deploy power at scale for the AI-driven data center era.

Bergen-Marelli Gensets are designed for continuous operation, typically running 7,000+ hours per year for 20+ years. They offer higher fuel efficiency and longer service intervals than high-speed engines, while delivering faster ramp response and greater operational flexibility than turbine-base systems. This balance of efficiency, durability and flexibility makes them ideally suited for base load and island-mode applications that can scale as AI power demand continues to grow.

## A Typical Entry-Level Configuration for an AI-Class Islanded Power Plant. 1x SHIELDX™: 1x 12MW Bergen–Marelli genset

This 1:1 ratio comfortably protects against typical 25–30% instantaneous load steps. If one generator drops off the bus unexpectedly, SHIELDX instantly injects active power to hold frequency and voltage steady while the remaining engines ramp-up or a standby unit synchronizes — with reserve capacity still available to manage transient events.


**Bergen + Piller ShieldX**

By increasing the SHIELDX ratio further, the system can cover up to 100% load steps, enabling the plant to ride through even extreme simultaneous events — such as a major AI-driven demand surge coinciding with multiple generator trips.

This level of plant-side protection, combined with on-site backup LNG storage, unlocks Tier IV-class resilience with five-nines availability, while allowing the engines to operate within their optimal ramp range for maximum efficiency.

With large-scale AI projects increasingly delayed by grid constraints, data center operators are turning toward behind-the-meter power generation. Langley's modular, built-to-stock architecture makes this approach readily scalable and ideal for phased expansions — from 50MW to over 500MW.

## Time to Power.

In today's market, the challenge isn't about favouring one technology over another — it's about delivering dependable power on time and at scale.

Grid connections for new AI-class data centres can take 6+ years to secure. Turbine-based islanded plants aren't far behind, often requiring 3 - 4 years to build and offering less agility in handling fast-ramping load profiles.

For many operators, an engine-based plant with SHIELDX™ stabilization offers a practical path: it can be deployed sooner, scaled in phases, and delivers the stability and resilience that AI workloads demand.

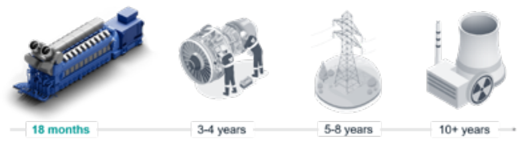


**Speed to power**

- A Bergen–Marelli power plant with SHIELDX™ stabilization can progress from contract signature to first batch ready for shipment in under 12 months, with Power-On in as little as 18 months.
- Once operational, the combined running cost per MWh — factoring Bergen's Long-Term Service Agreement (LTSA) and fuel — can be competitive with, and in some regions lower than, prevailing grid-power prices, while delivering superior stability, reliability, and control over future energy costs.

## The holistic result.

Data-center-side innovations have made significant progress in moderating AI-driven power demand profiles. They shrink the size and speed of the ramps that reach the upstream supply, but they cannot cap the sharpest peaks and they do not shield an islanded plant from sub-second, multi-megawatt load steps.

True stability has to start at the source — at the power plant. That first layer, the foundation of a resilient island-mode system, must come from plant-level dynamic stabilization.

Building your power plant around SHIELDX™ means:

- Engines run steady and efficiently — with fewer starts, less wear, lower emissions, and lower OpEx.
- UPS and IT loads receive premium-quality power — as if tied to a strong public grid.
- Operators can right-size, not over-size the power plant avoiding the fuel, cost and emission penalties of keeping extra capacity online.

The future of behind-the-meter AI power can be a hybrid form, combining SHIELDX stabilization with AI-ready UPS, rack level smoothing and raised power floors, or delivered as a standalone, total plant-level resilient solution.

Together these approaches enable a leaner and more reliable island-mode power plant at scale

For further information email **PowerOn@langleyholdings.com**