ON energy™

March 2026

AI Needs Power. Here's the Opportunity.

Alan Cooper
CEO & Co-founder, ON.energy

Start with what's at stake. The United States is in a race to build the AI infrastructure the world will run on. Not a product cycle. Not a software update. The next industrial revolution.

That race is real. And it runs on electricity. And the countries that build the necessary infrastructure will shape the global economy for decades. Every data center built, every Graphics Processing Unit (GPU) cluster commissioned, every model trained is the result of electrons moving through silicon. Power is not a supporting actor in the AI race. It’s not even the main character. It’s the stadium in which the race is being held.

How you build matters as much as how fast you build. If AI infrastructure destabilizes the grid, crowds out clean energy, and drives up electricity bills for ordinary consumers, the public will push back. The buildout will slow. The window will close.

What makes AI power different?

Ripple.

Most industrial power loads are predictable. A factory runs a shift. An office building peaks in the mid morning and mid afternoon. Grid operators have spent a century learning to plan for these patterns.

AI training is different. Thousands of GPUs are grouped together to build models & process information in unison. These clusters don't draw power steadily. They operate in phases. High-intensity computation during active training. Checkpointing. Data loading. Cross-cluster synchronization. This batch processing causes large, rapid step changes in power demand: hundreds of megawatts, ramping up and down about every 10 seconds.

From the grid's perspective, a large AI training facility looks like a city that keeps going dark and lighting up again. Hundreds of times a day. That creates a problem. Call it the AI Training Ripple.

The power oscillations generated by GPU workloads don't stay inside the building. They propagate back through the point of interconnection into the transmission system. They stress grid equipment. They degrade power quality for every other customer on the same system. In facilities that co-locate gas turbines or backup generation, the Ripple creates mechanical torque oscillations in turbine shafts — accelerating metal fatigue, shortening equipment life, forcing generators to run outside their efficiency band.

The workload paying the bills is quietly damaging the machines keeping the lights on.

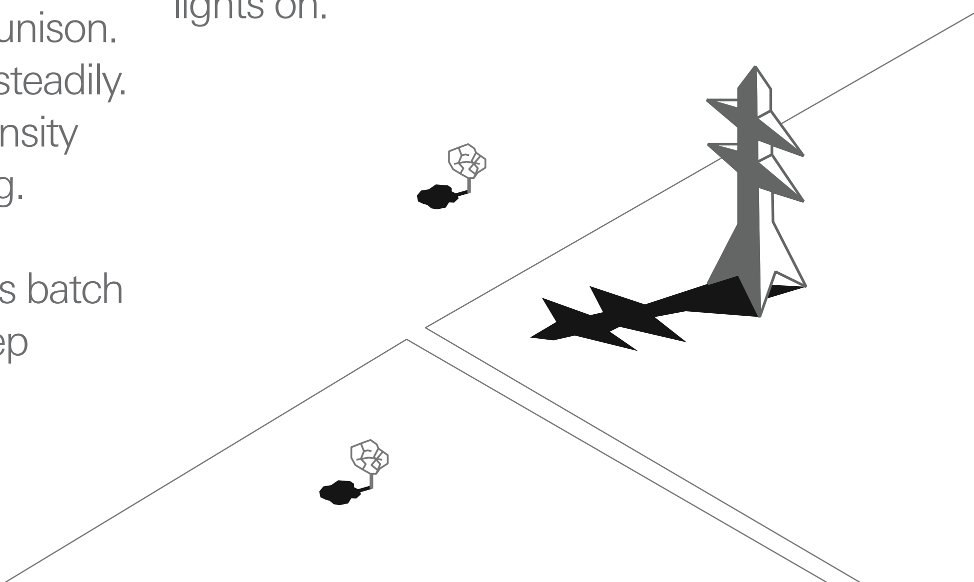

Size.

Data centers are big because AI demands it. Training a single large model can require hundreds of thousands of GPUs running for months. That compute draws power at a scale commercial real estate has never seen, and the facilities are still growing. Five years ago, a large campus drew 100 megawatts. Today, hyperscalers design for a gigawatt at a time.

That scale changes what these facilities are. A 100-megawatt data center is a large industrial customer. A gigawatt campus is a transmission-level resource. When it connects, the grid feels it. When it disconnects, the grid feels that too. The grid runs on balance. Remove a gigawatt of demand in milliseconds and the system has to rebalance instantly or frequency drifts, cascades start, and the lights go out somewhere else.

The disconnect is the problem. Data centers are designed to protect themselves. The moment they sense a grid disturbance, they drop off — automatically, in milliseconds, no operator in the loop. Right instinct for protecting billions in compute. Wrong behavior for the grid.

In July 2024, a lightning strike in Virginia triggered that instinct across 60 data centers simultaneously. PJM absorbed a sudden 1,500-megawatt loss of load. It worked…barely. Ten months later, a fault in Ireland shed 387 megawatts in an instant, more than half the country's total data center demand, forcing emergency intervention. That same month, Spain and Portugal went dark for nearly 24 hours—Europe's worst blackout in two decades. Different cause, same physics: a large, fast imbalance the system couldn't absorb.

The facilities that tripped in Virginia and Ireland were protecting themselves. A gigawatt campus doing the same thing is a different order of problem entirely.

Interconnection.

The interconnection queue wasn't built for this.

Getting a large load connected to the grid has always been slow. The Ripple and the Size problem make it slower. A gigawatt campus with oscillating, unpredictable load behavior requires grid impact analysis that didn't exist five years ago. In PJM, interconnection has historically taken close to eight years. Reforms underway aim to cut that to one to two years. Neither number works for a data center that needs to be online in 18 months.

Most developers go around it. Build your own power. Skip the queue. Rational at the facility level. At the system level, it solves nothing.

The federal government has drawn the line. The White House Ratepayer Protection Pledge, issued March 4th, 2026, requires hyperscalers to build, bring, or buy their own power, pay for all new delivery infrastructure, contribute backup generation during scarcity events, and ensure none of these costs reach ordinary ratepayers.

The mandate is real.
The question is no longer whether to solve these problems,
it is how.

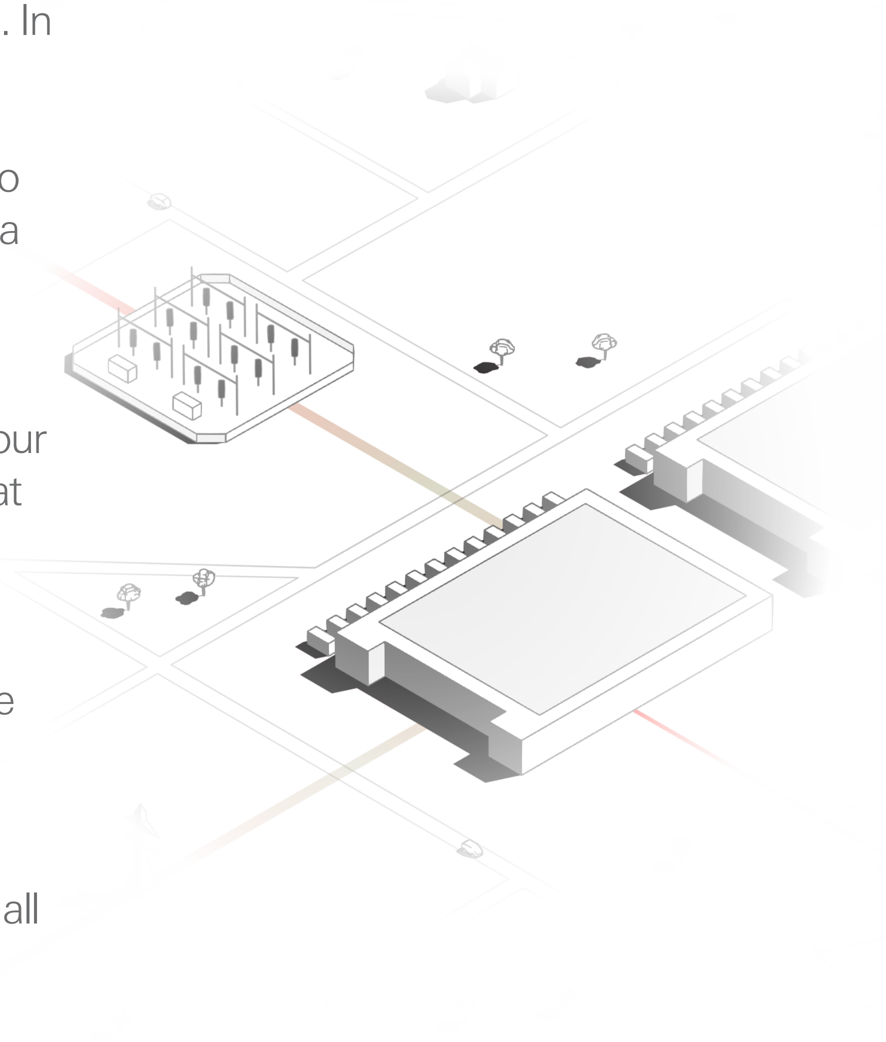

What do data centers need?

Protection from the Grid.

Every Uninterruptible Power Supply (UPS) deployed in a data center today protects the compute. Almost none protect the cooling.

For most of the industry's history, that was fine. Legacy servers had enough thermal tolerance that a brief interruption to cooling was manageable. Modern GPU clusters do not.

GPU clusters run at heat densities that liquid cooling is barely keeping pace with under normal conditions. Take the cooling offline and temperatures don't drift up, they spike. A UPS that keeps servers drawing full power without the cooling to match isn't buying time. It's accelerating the problem.

Protection for the Grid.

The problem isn't what data centers consume. It's how they behave when things go wrong.

When transmission faults hit, facilities switch to backup automatically, shedding load at a speed and scale the grid wasn't designed to absorb. A gigawatt campus isn't a grid customer. It's a grid participant, whether it chooses to be or not. At that scale, disconnecting isn't self-protection. It's a grid event.

The architecture to behave differently—demand response, frequency regulation, ancillary services—exists. Most facilities just can't participate. Not because the programs aren't available, but because responding requires modulating load on command. And the only levers most operators have are the ones that also put compute at risk. Shed load and you risk dropping servers. Cut power and you risk the cooling. There's no safe middle. So operators opt out entirely, and the grid loses a resource it increasingly needs.

That gap is a problem ON's AI UPS closes.

One Architecture. Both Problems. A Shameless Plug.

ON's medium voltage AI UPS™ sits at the campus medium-voltage bus, at the grid interconnection point, upstream of everything. It acts as a firewall between the grid and the data center. Protecting both from each other. More value than a traditional low voltage UPS, all the value of energy storage, rolled into one.

This architecture makes whole-campus protection possible without replicating UPS infrastructure hall by hall. Compute and cooling protected together. And because it sits upstream of both, it suppresses the GPU load ripple at the source, before it stresses generation equipment, before it reaches the transmission system.

Simplified electrical designs. Data halls that carry substantially more IT load per building, at lower cost, with better resilience.

And because it acts as a medium voltage UPS coupled with short-duration storage, it becomes the hardware foundation for grid participation. Frequency response. Voltage support. Demand flexibility. One decision, one system, in both directions: downstream through the campus, and upstream into the grid.

A facility that can demonstrate ride-through and offer credible grid services is a different interconnection applicant. It is absorbing risk the utility was otherwise being asked to carry. That changes what the safety review concludes. That changes the timeline.

That is speed to power.

Proof.

In February, ON.energy completed live validation of the AI UPS™ at the National Laboratory of the Rockies, a U.S. Department of Energy facility and the only place in the Western Hemisphere with both a grid simulator and a data center simulator.

We introduced load swings from 30% to 100% of capacity in intervals as short as 10 milliseconds. The upstream grid profile remained flat. The system absorbed the transients in real time, maintaining a controlled ramp rate and stable state of charge throughout. AI load volatility can be contained before it reaches the grid. That is not a claim. It is a measured result.

The first and only solution to not only meet, but exceed ERCOT voltage ride-through requirements for Large Loads.

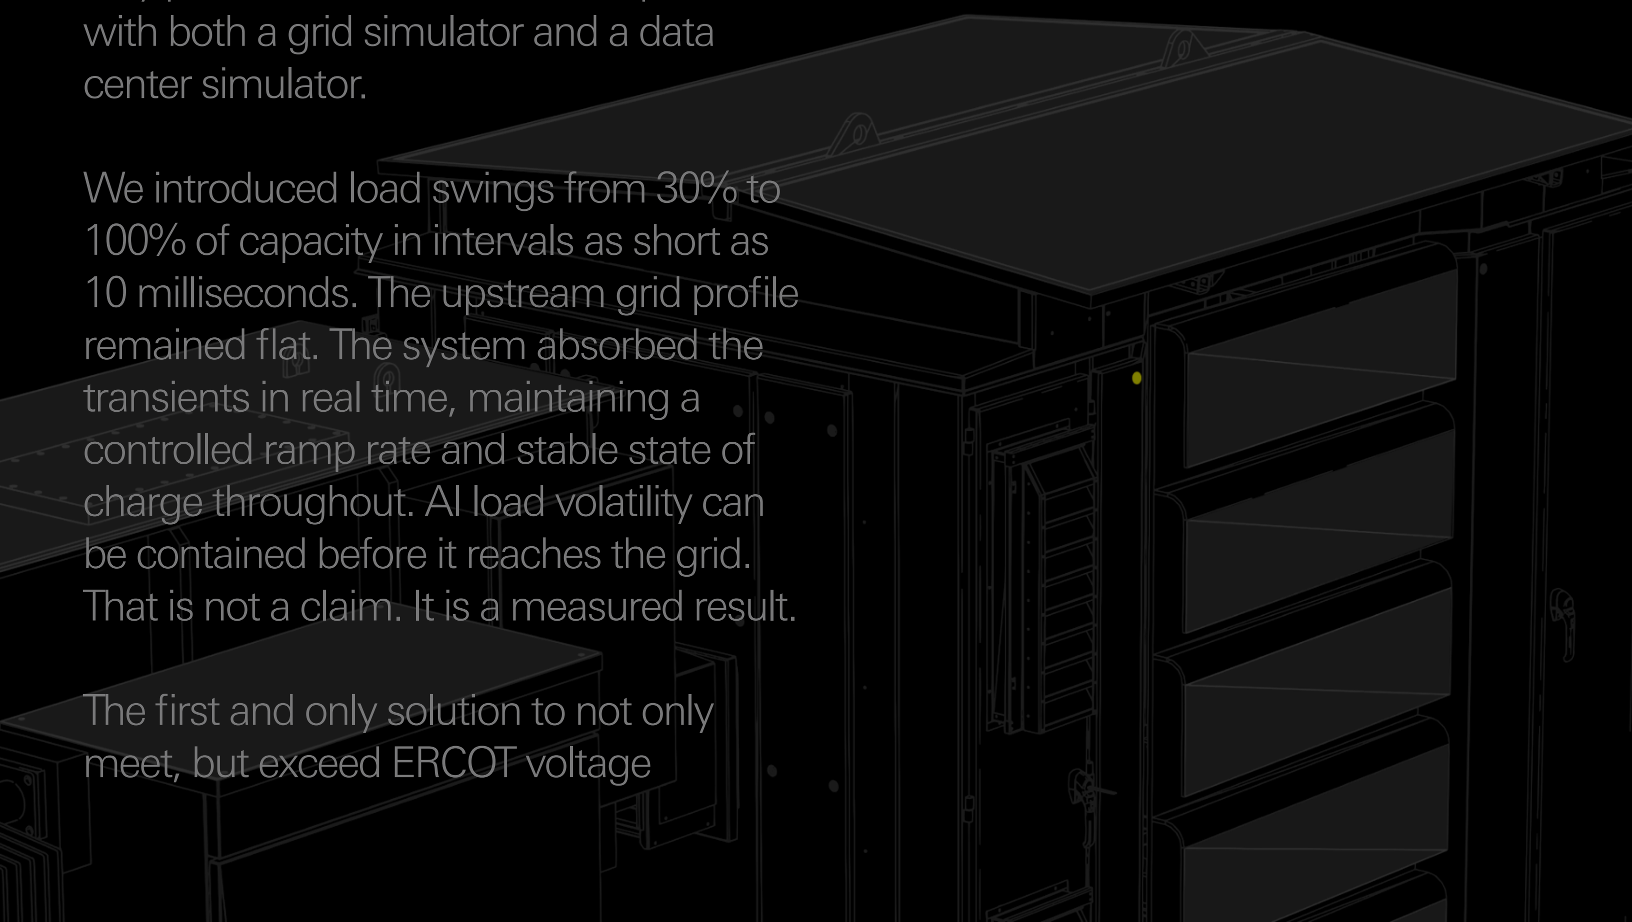

ON.energy is actively deploying over 2 GW of AI UPS across grid-tied and fully off-grid data center projects. Another 3 GW has been selected for data center projects across the US.

Infrastructure designed for what AI actually demands: ride-through at medium voltage, with power and cooling protected together from the point of interconnection.

One system. Grid-safe.

Why We Are Here.

I have spent ten years working on tough power problems that don't make headlines. Behind-the-meter resiliency systems for large energy users. Distributed grid-tied storage that keeps the grid balanced when supply and demand pull in different directions. Now that work is at the center of the most consequential infrastructure buildout in American history.

The AI buildout is the largest coordinated energy deployment America has ever attempted, happening right now, at a pace nobody anticipated. The opportunity is to build power infrastructure at the data center that serves the grid, not strains it.

To use the largest movement of capital in human history to pull forward the grid we actually need.
What the people building these facilities need is energy architecture that is simple enough to adopt, incentivized enough to choose, and robust enough to trust. Energy is not their core business. It shouldn't have to be.

That is the work. Not just building the AI UPS, but establishing the standard for how data centers connect to the grid. Working with developers, utilities, regulators, and grid operators to make grid-safe architecture the default rather than the exception.

We are not going to rebuild the grid for AI. We are going to make AI infrastructure a support system for the grid. That is the standard we are here to set.
That is the opportunity.

https://on.energy