

247 ENERGY

Grid Congestion and the Bridge Past the Queue

How on-site storage and rapid containerised power let industrial sites grow when grid capacity has run out.

James Troch

Chief Executive Officer

247 Energy NV

June 2026

The Constraint Has Moved

For most of the last century the limit on industrial growth was capital or demand. A company that could fund a new line and sell what it made could build. That is no longer reliably true. In a growing number of regions the binding constraint is the grid itself. A business can hold the site, the capital, the orders, and the permits, and still be told there is no capacity to connect, and none expected for years.

I have watched this shift move from an occasional frustration to a structural feature of the markets we work in. Connection queues that once cleared in months now stretch beyond the planning horizon of the projects waiting in them. The grid was built for a world of large, controllable power stations and steady demand. It is now asked to carry electrified transport, electrified heat, new digital load, and reindustrialisation at the same time, and in many places it has simply run out of room.

This paper is not written to sell a product. It is written to explain a problem that is reshaping where and how businesses can expand, and to set out the two practical answers available today. The first is to change the shape of the demand a site presents to the grid, so that a smaller connection carries a larger operation. The second is to generate power on site, quickly, where the grid cannot deliver it at all. Neither requires waiting for the network to catch up.

The companies that treat grid access as a fixed external constraint will wait. The companies that treat it as a problem they can engineer around will grow. I believe that distinction will separate the winners from the rest over the next decade, and the pages that follow explain why.

James Troch, Chief Executive Officer, 247 Energy

A Grid That Ran Out of Room

How congestion became the default condition

Grid congestion happens when the physical network cannot move all the power that generators want to inject or that consumers want to draw. It is not a shortage of electricity. It is a shortage of capacity in the wires, transformers, and substations that carry electricity from where it is made to where it is used. When that capacity is exhausted, the operator has two choices: refuse new connections, or curtail existing flows. Across much of northwestern Europe, and increasingly elsewhere, both are now routine.

The causes are structural rather than temporary. Renewable generation has been added far faster than the network that must carry it, and it often connects in places that were never designed to host power production, such as rural land with good wind or sun but thin grid infrastructure. At the same time demand is rising and concentrating. Electric vehicle fleets, heat pumps, and large digital facilities all draw heavily, and they tend to cluster. The result is a network stressed at both ends, with generation straining to get on and demand straining to get more.

Building new grid capacity is the obvious answer and the slowest one. A major reinforcement, a new line or a rebuilt substation, can take the better part of a decade once planning, permitting, land access, and construction are accounted for. Demand does not wait that long. A logistics operator electrifying a fleet, a manufacturer adding a line, or a developer building out a site needs power on a timescale of months. The mismatch between how fast demand moves and how slowly the grid can grow sits at the heart of the problem.

Grid congestion is not a shortage of electricity. It is a shortage of room to move it, and that room cannot be built fast enough to meet demand on its own.

For the business on the receiving end, the experience is blunt. A connection request returns a date years out, or a quoted capacity well below what the operation needs, or a reinforcement cost so large it changes the economics of the whole project. Expansion plans stall, not because the market is absent or the capital is unavailable, but because the single input that everything else depends on cannot be secured. This is a new kind of constraint, and most planning processes were not built to anticipate it.

What makes the situation manageable is that the grid connection is not the only way to get power, and the size of a connection is not fixed by the size of an operation. Both assumptions can be challenged with technology that exists and is deployed today. The rest of this paper explains how.

Reshaping the Load

Why a smaller connection can carry a larger operation

The capacity of a grid connection is sized to the peak demand it must serve, not the average. A site that draws two megawatts for fifteen minutes a day and far less the rest of the time still needs a connection rated for that brief peak. Most of that expensive capacity sits idle most of the time. Energy storage attacks exactly this inefficiency. By charging when demand is low and discharging to cover the peaks, an on-site battery lets a site present a flatter, lower profile to the grid than its actual consumption would suggest.

The practical consequence is that storage decouples what an operation uses from what its connection must be rated to deliver. A site facing a connection limit can install storage to shave its peaks below that limit and continue to expand the work it does behind the meter. The grid sees a steady, predictable draw. The operation behind it can grow, add equipment, and run harder, because the spikes that would otherwise breach the connection are absorbed and released locally.

The size of a grid connection is set by the peak, not the average. Storage lets a site grow the operation without growing the peak the grid has to serve.

This same mechanism does more than relieve a capacity limit. Where networks charge for peak demand, flattening the profile lowers those charges directly. Where supply is intermittent or the connection is weak, storage holds energy through the gaps and steadies the supply the site actually experiences. A single asset addresses a capacity constraint, a cost problem, and a reliability problem at once, which is part of why it has moved so quickly from a niche measure to a mainstream tool for energy-intensive sites.

The technology matters here, because storage installed close to people and stock carries requirements that go beyond raw performance. It must be safe to site near occupied buildings, quiet enough not to create a nuisance, and durable enough to cycle hard every day for years without rapid degradation. Storage built on supercapacitor technology answers these directly. It does not suffer thermal runaway, the chain reaction that drives battery fires, because the failure mode is physically absent rather than merely managed. It tolerates frequent deep cycling without the wear that shortens the life of conventional cells. And it switches between charge and discharge instantly, which matters when it is standing in for a connection that must never be late.

Sizing storage to a site is an engineering exercise, not a guess. The load profile, the peak pattern, the connection limit, and the growth plan together determine how much storage is needed and how it should be operated. Done well, the result is a connection that no longer caps the business, achieved in the time it takes to install equipment rather than the time it takes to rebuild a substation.

Power That Arrives in Hours

On-site generation where the grid cannot deliver at all

Reshaping demand works when a connection exists but is too small. It does not help when there is no usable connection at all, or when a site needs more power than any plausible local upgrade can provide in the available time. Construction sites, temporary operations, remote facilities, and sites waiting years for a connection all face this harder version of the problem. For them the question is not how to use a connection more cleverly but how to generate power independently, cleanly, and fast.

Containerised power generation answers that question. A generation unit delivered in a standard container can be transported to a site and brought into operation in under four hours. It needs no permanent grid infrastructure and no long construction programme. For an operation that would otherwise be stalled for years waiting in a connection queue, the difference between four hours and four years is the difference between proceeding and not.

The fuel matters as much as the speed. Diesel generation, the traditional fallback, is loud, dirty, and expensive to run at scale. Generation fuelled by liquefied natural gas changes that balance. It produces roughly a quarter less carbon dioxide than diesel for the same output, and it cuts nitrogen oxide emissions, the pollutants most associated with local air quality and health, by close to ninety-nine percent. Over the life of a deployment the total cost of ownership can be around half that of an equivalent diesel installation. For a bridging solution that may run for months or years while a connection is awaited, those differences compound into a materially better outcome on cost, emissions, and local acceptability.

Containerised generation is not a permanent replacement for a grid connection, and it is not presented as one. It is a bridge. It lets a site operate, earn, and build during the years it would otherwise lose to the queue, and it can be scaled down or removed as grid capacity finally arrives. Treated that way, as a temporary instrument deployed deliberately rather than a permanent dependence, it converts a multi-year delay into a manageable interim.

The same logic extends to resilience. A site with independent generation available is not exposed to the full consequence of a grid outage or a curtailment instruction. The capability installed to bridge a connection gap doubles as insurance against interruption, which for many industrial and logistics operations is worth as much as the capacity itself.

Designing Around the Constraint

Hybrid systems and staged expansion

The two answers are stronger together than apart. Storage reshapes the demand a connection must serve. On-site generation provides power the connection cannot. A site can combine them: storage to flatten peaks and ride through short gaps, generation to cover sustained shortfalls or to operate ahead of any connection at all. Designed as one system rather than two separate purchases, the combination can match almost any constraint a site faces, from a connection that is merely too small to one that does not yet exist.

Staging matters. A growing operation rarely needs its full future capacity on day one. A system can begin with storage against an existing connection, add generation as load grows beyond what the connection can carry, and scale back the generation as grid reinforcement eventually lands. This turns a binary decision, wait for the grid or do not build, into a sequence of smaller, reversible steps that track the actual growth of the business rather than betting everything on a single connection date.

Grid access stops being a fixed external limit and becomes a variable the business can engineer, stage, and control on its own timetable.

This reframing is the real shift. Treated as a fixed external limit, a connection constraint forces a business to wait on a process it does not control. Treated as an engineering problem, it becomes a set of choices the business can make and revise: how much storage, how much generation, in what sequence, operated under what logic. The constraint does not disappear, but it stops dictating the timetable.

Designing these systems well requires seeing the site as a whole. The load profile, the growth plan, the connection status, the local constraints on noise and emissions, and the resilience requirements all interact. A peak-shaving battery sized without reference to the growth plan will be wrong within a year. Generation specified without regard to local air quality limits may not be permissible. The value comes from designing the energy system and the business plan together, so that each supports the other.

None of this is speculative. The components are in service today, sized and operated for real sites with real constraints. What is still uncommon is the habit of treating grid access as something to be engineered rather than awaited. That habit, more than any single piece of equipment, is what separates the sites that keep growing from the sites that stall.

A Problem That Will Spread

Why congestion is structural, not a passing phase

It would be comfortable to treat grid congestion as a temporary growing pain that reinforcement will soon resolve. The evidence points the other way. The forces driving congestion are the electrification of transport and heat, the build-out of renewable generation, the arrival of large new digital loads, and the reindustrialisation of advanced economies. All of them are accelerating, and all of them press on the same finite network at the same time.

Grid reinforcement is not accelerating at a matching pace. It is constrained by permitting timelines, supply chains for transformers and cables, skilled labour, and public opposition to new infrastructure. Even where investment is rising, the lead times are measured in years and the backlog is already large. The gap between the speed of demand and the speed of supply is widening, not closing, in most of the markets where it has appeared.

The geographic spread is widening too. Congestion appeared first in regions with high renewable penetration and dense demand, but the same pattern is emerging wherever electrification and new load growth outpace the local network. A constraint that was once a regional peculiarity is becoming a general condition of operating an energy-intensive business in a developed economy. Planning on the assumption that it will pass is planning to be caught out.

For businesses, the implication is straightforward. The ability to operate without full reliance on a timely grid connection is shifting from a niche advantage to a core competence. The firms that build that capability early will expand into the space the constraint creates. The firms that wait for the grid will find that waiting is itself a competitive decision, and usually the wrong one.

THE STRATEGIC CASE

Turning a Constraint Into an Advantage

What early movers gain

A constraint that affects everyone equally is not a disadvantage. It is a sorting mechanism. When grid access limits an entire region, the businesses that can operate around it gain ground on the ones that cannot. The early mover does not merely avoid the delay. It expands into demand that competitors are unable to serve, secures sites that would otherwise be unviable, and builds operational know-how that compounds with each project.

The financial logic favours action over delay in most cases. The cost of installing storage or bridging generation is known, bounded, and incurred against a clear return: the revenue from an operation that can run now rather than years from now. The cost of waiting is harder to see but often larger. It includes foregone output, lost market position, and the option value of a site that sits idle. Framed honestly against the cost of doing nothing, on-site energy infrastructure is frequently the cheaper choice, not the more expensive one.

There is also a resilience dividend that does not appear on the connection paperwork. A site engineered to operate independently of a timely grid connection is, by the same design, more robust against outages, curtailment, and price spikes. The capability bought to solve a capacity problem keeps paying as insurance against the volatility that increasingly characterises power supply. For operations where downtime is expensive, that secondary benefit can justify the investment on its own.

The strategic question for a business facing a connection constraint is therefore not only how to get connected, but whether full dependence on the connection is the right design at all. For a growing number of operations, the better answer is to build a degree of energy independence into the site from the outset, and to treat the grid as one input among several rather than the single point of failure on which everything rests.

THE TIMING

Why the Window Is Now

Timing favours acting before the constraint fully arrives rather than after. The sites that secure storage and on-site generation while equipment, installers, and connection slots are still available will be better placed than those that move once congestion has become acute and everyone is competing for the same solutions at once. A constraint visible on the horizon is cheaper to address than a constraint that has already bitten.

The supply side reinforces this. The same pressures that congest the grid, namely rising demand for transformers, cables, skilled installers, and equipment, also tighten the market for the on-site alternatives. The cost and lead time of a storage or generation deployment are lowest when demand for them is lowest, which is now, before the wave of constrained sites all reach for the same answer in the same window.

Regulatory direction adds to the case. Across most jurisdictions the policy trajectory is toward integrating storage and distributed generation more deeply into how the system is planned and rewarded, not less.

Acting early positions a site to benefit from that direction rather than to retrofit against it later, and to shape its energy design while the rules are still settling rather than after they have hardened.

The core point is simple. Grid congestion is a structural, spreading, and durable constraint, and it can be engineered around with technology that is available and proven today. The advantage goes to the businesses that recognise this early and design for it, not to those that wait for a network to catch up with a demand it was never built to carry. The window to move ahead of the constraint is open now, and it will not stay open indefinitely.

2 4 7 ENERGY

Building the Energy Layer for Constrained Sites

247 Energy develops, builds, and operates distributed energy solutions for sites where the grid alone is no longer enough. Our work spans three areas that map directly onto the problem this paper describes: energy storage sized to reshape the demand a site presents to the grid, containerised power generation that brings clean capacity to a site in hours, and the development of utility-scale battery storage parks that add stabilising capacity to the networks themselves.

Our storage is built on supercapacitor technology, chosen for exactly the conditions on-site deployment demands. It does not suffer thermal runaway, so it can be sited safely close to people and stock. It tolerates frequent deep cycling without rapid degradation, so it earns its keep over a long life rather than wearing out under heavy use. It runs quietly, at around fifty-five decibels, and it switches between charge and discharge instantly, which matters when it is standing in for a connection that cannot be late.

Our containerised generation runs on liquefied natural gas and is operational in under four hours from delivery. Against an equivalent diesel installation it produces roughly a quarter less carbon dioxide, cuts nitrogen oxide emissions by close to ninety-nine percent, and can roughly halve the total cost of ownership over the life of a deployment. As a bridge past a connection queue, or as resilient backup, it lets a site operate during the years it would otherwise lose.

At utility scale, we develop, build, and co-invest in battery energy storage parks rated at 100 MW and above. Our current European project pipeline spans two regions and five countries, with a combined capacity of 505 MW / 2,025 MWh. We retain skin in the game in the assets we develop, because we believe long-term performance is where the real value of this infrastructure is captured, and we work with co-investors and developers who share that long-term view.

For businesses facing a connection constraint, and for partners interested in co-investment in storage and on-site power, we welcome direct engagement. We can discuss specific situations in detail, from sizing an on-site system around a constrained connection to participating in utility-scale projects, under appropriate confidentiality arrangements. Our aim is to work with parties who understand the constraint, share our conviction about how to engineer around it, and take a long-term view of the value it represents.

247 Energy NV

Schaarbeekstraat 20E/11 | 9120 Beveren, Belgium

+32 3331 0000 | storage@247.energy | 247.energy

Copyright 2026 247 Energy NV. All rights reserved. This paper is intended for informational purposes only and does not constitute an offer or solicitation of any investment, financial product, or service. Statements regarding future events or outcomes are forward-looking and subject to risk and uncertainty.