

The Constraint Stack

Why Distributed Power Projects Fail Before They Start

By Chad Forbes, Aurelia Technologies

Executive Summary

Many distributed power projects do not break down in procurement, construction, or operations. They often do so earlier – in how the problem is framed and due to a simplistic, truncated analysis leading to value destruction later.

When a project is treated as a straightforward equipment purchase or a line-item capital exercise, it is often filtered through assumptions too narrow for the conditions it must actually satisfy. Sometimes the result is visible: delay, redesign, rework, or abandonment. Often, however, the project proceeds – but in a form that leaves value unrealized because the analysis never fully accounted for the real constraints shaping the outcome.

Distributed power decisions are often evaluated through a small set of obvious variables. These variables are easier to model, compare, approve, or defend internally. In practice, though, project success is governed by a broader **constraint stack**: the interacting timing, operational, site, commercial, and organizational realities that collectively determine whether a solution can deliver.

As constraints accumulate, the economics change. Substitutes change. Decision criteria change. In a market defined by volatile lead times and shifting grid reliability, the most valuable attribute of a power project is its **optionality** – the ability to remain viable as project conditions evolve.

The issue is not a lack of technical options; it is a mismatch between project complexity and the simplicity of the evaluation used to define it. This paper proposes a more practical approach to screening:

Start with a better first question:

What conditions must this project satisfy to succeed – and which constraints will determine whether the answer works in the real world?

Many Power Decisions Lose Value Before Any Equipment Is Chosen

Power projects can lose value in framing before a single piece of equipment is selected.

The most expensive mistake in distributed power is usually analytical, not technical. A project may be defined too narrowly, evaluated too early through a familiar category, or pushed through a decision path easier to manage internally than representative of the actual challenge. Sometimes the result is obvious: the project stalls or is delayed. Often, the project moves forward – but with an architecture, structure, or operating assumption that leaves performance, resilience, flexibility, or economic value on the table.¹

This pattern matters more today because conditions surrounding distributed power are harder, not easier, to simplify. Load growth, constrained utility timing, site expansion pressure, equipment lead times, and rising operational expectations create situations in which “wait for utility power” or “buy standard backup” no longer captures the full decision. Power availability itself has become a gating condition for readiness,

¹ LBNL (2024), *Queued Up*.

commissioning schedules, and site viability, while interconnection delays and infrastructure bottlenecks² increasingly shape timing and form.³ This is no longer just an issue for large data centers; many utilities have implemented Large Load Integration (LLI) or similar 'Fast-Track' protocols that trigger long study cycles for even significant industrial expansions. Similarly, congestion and increased grid demand have increased scrutiny and evaluation times for smaller, distribution side capacity increases or new interconnections. These protocols mean that for the project owner, the utility is no longer a service provider to be managed, but a binding constraint to be solved.⁴

The central question is often framed too late or too loosely. Teams frequently begin with a familiar procurement or engineering question: *Which solution should we buy?* In more constrained environments, the more useful question is broader: *What is actually stopping this project from succeeding on time, on budget, and in a form operations can live with?* For some sites, the answer remains straightforward. For others, the decision is shaped by timing pressure, uncertain load growth, deployment constraints, reliability exposure, commercial limitations, or internal approval friction – often in combination. Once several of these conditions are active, the project stops behaving like a simple equipment purchase and begins behaving like a **constraint problem**. Many projects drift there long before anyone recognizes it.

The Real Challenge Is Usually a Stack, Not a Single Variable

Most distributed power opportunities are shaped by multiple interacting concerns.

Stakeholders are rarely looking at the “wrong” variable. More often, under decision pressure, they gravitate toward variables easiest to compare, defend, or move through approval channels: price, efficiency, utility timing, and reliability often serve as proxies for a more complex set of conditions. This is understandable. It is also where distinctions can get compressed out of the analysis.

A site may not simply need “power.” It may need power before utility service is available, on a footprint that limits conventional deployment, in a form that can support staged or uncertain growth, under internal capital or commercial restrictions, and with sufficient confidence for approval and implementation.

Each condition matters individually. What changes the character of the decision is their interaction. A project that is delayed is one thing. A project that is delayed, space-constrained, operationally sensitive, and commercially inflexible is another.

The Architecture of a Constraint Stack

In practice, these conditions cluster into recurring families:

- **Time and access:** focus on speed-to-power and the physical reality of deployment.
- **Economic and commercial:** focus on how the project is financed, justified, and absorbed.
- **Operational and performance:** focus on the site’s actual load/uptime needs.
- **Site and execution:** focus on permitting, footprint, and the “built” reality.

These families represent the “binding” realities of modern power projects. For example, **Time and Access** has become a primary gating factor in many markets as interconnection wait times are extended.⁵ Similarly,

² Wilson, S&P Global (2025), *Grid Congestion & Load Growth*.

³ IEA (2025), *Electricity 2025*; IEA (2025), *Mid-Year Update*.

⁴ ERCOT (2026), *Large Load Integration*.

⁵ LBNL (2024).




Figure 1. Most distributed power projects are not governed by one variable. They are shaped by a set of interacting operational, commercial, and deployment constraints.

Economic viability is increasingly dictated by the "cost of friction" – where grid congestion forces a shift toward behind-the-meter solutions.⁶ Even when these are solved, **Operational** success remains fragile; research across mission-critical infrastructure operations consistently finds that the majority of unplanned downtime traces to process, management, and planning failures rather than equipment malfunction — a pattern clearly documented in high-availability environments.⁷

Individually manageable, in combination they change the screening logic. A project can look clean when each issue is isolated. Once several interact, the analysis must account for that complexity.

None of these situations are inherently unsolvable. They do, however, require evaluation logic that reflects reality rather than a simplified category.

Why Sophisticated Organizations Still Compress the Analysis

Weak distributed power decisions rarely come from unsophisticated teams. They arise in organizations making difficult, multi-variable decisions under real-world pressure. The reason this matters is not academic; compressed analysis has compounding economic consequences.

⁶ Wilson, S&P Global (2025).

⁷ Uptime Institute (2025), *Annual Outage Analysis: Metrics and Trends*

Opportunities still enter the system through familiar lanes: backup purchases, utility timing issues, facility upgrades, procurement events, or capital comparisons. These lanes exist for good reasons; they help organizations move. The problem arises when the project no longer fits neatly inside them.

Common compression patterns include, for example:

- **The Optimism Bias:** assuming utility interconnection queues will shrink or can be “managed” through escalation.⁸
- **Categorical Narrowing:** reducing the decision too early to capital cost or standard hardware comparisons.
- **Architectural Passivity:** treating reliability as a simple specification rather than an architectural design question.
- **The “Least Regret” Trap:** defaulting to the easiest option for internal approval, which often results in the highest operational friction.

Project Situation	Typical Constraint Collision	Why Simplistic Screening Fails
Utility Power is Too Slow	Interconnection delays, critical commissioning timelines, load ramp-up pressure	A simple “wait for utility” assumption underestimates schedule risk and forces temporary solutions that may not scale or protect uptime.
The Site is Temporary, Staged, or Uncertain	Short-term leases, phased commissioning, evolving footprint	Screening as a standard backup or permanent installation misses flexibility needs; the “good enough” solution today can block tomorrow’s growth.
Capacity Needs Will Grow in Steps	Load growth uncertainty, phased expansion budgets, future optionality	Early sizing or cost comparisons ignore staged growth; project appear simpler than they are, creating retrofits or costly adjustments later.
The Load Will Not Behave Neatly	High variability, unpredictable duty cycles, operational constraints	Treating load as uniform or predictable hides the need for flexibility, redundancy, and operational support, underestimating operational risk.
The Site Itself is Constraining the Answer	Footprint, access, environmental limits, permitting	Standard solutions are forced into inadequate spaces; early screening overlooks modular or creative architectures, reducing both performance and optionality.

Projects often lose quality quietly. Some are screened out because the process was designed for a different decision type. Others are forced into a convenient answer because it is easier to approve.

When assumptions carry too much weight – e.g., “utility timing will improve,” “temporary solution is sufficient,” or “we can size later” – they may obscure the real constraints. A process that cannot distinguish defensible from convenient assumptions will struggle in more constrained environments.

The struggle becomes especially expensive near category boundaries: when something that initially looks like a backup or temporary power decision is actually critical for uptime, phased expansion, or future flexibility. At that point, the decision is no longer just technical; it is economic and organizational.

⁸ IEA (2025), *Mid-Year Update*.

Once Constraints Stack, the Economics Change

One of the most common mistakes in distributed power evaluation is assuming the economics are obvious. Often, they are not.

As constraints stack, value stops behaving like a simple purchase-price question and starts behaving like a project economics question. Financial consequences no longer lie only in equipment cost. They appear across schedules, operations, risk exposure, and future flexibility.

A distributed power decision can create or destroy value through four broad levers:

- Profit/Productivity
- Cost/Efficiency
- Risk/Resilience
- Strategic Position



Figure 2. Most projects create tradeoffs across several of these value dimensions at once.

A delay in power availability affects more than capex timing. It can disrupt operations, idle labor and equipment, derail construction sequencing, shift revenue timing, tie up capital, and trigger temporary workarounds.

Likewise, an operating architecture that looks acceptable in a narrow comparison can create costs elsewhere – through uptime exposure, support burden, recoverability risk, or the inability to adapt cleanly as the site evolves.

Optionality deserves explicit treatment. Preserving future pathways – through phased deployment, modular expansion, commercial flexibility, or avoiding premature lock-in – is often

CASE IN POINT:

FROM CONTINGENCY TO OPERATIONS

Illustrative composite derived from recurring distributed power deployment patterns.

The Context: An industrial facility deployed a 2 MW N+1 gas generation system to improve resilience against infrequent utility outages. The project was evaluated primarily through risk reduction and internal operational confidence: utility service would remain the primary operating source, while onsite generation would function as contingent infrastructure during rare interruptions or maintenance events. Expected runtime was 50-100 hours per year, leading the project to prioritize upfront capital cost and standby-oriented operating assumptions.

The Friction: Over time, the facility's operating environment evolved in ways that were not fully reflected originally. Although prolonged outages remained relatively uncommon, power quality disruptions, utility switching events, peak-demand exposure, and curtailment risk began creating disparate operational impacts. As continuity sensitivity increased, even brief power quality events risked production instability, restart activity, scheduling disruption, or material loss.

In response, operators gradually expanded the role of the generation system. The units were increasingly dispatched during forecast grid stress periods, sensitive production windows, and selected peak-demand intervals to reduce operational exposure and avoid disruption costs. Runtime expanded severalfold beyond the original standby assumptions, materially increasing fuel consumption, maintenance requirements, compliance obligations, and operator involvement.

The Takeaway: The project itself did not fail technically. The issue was that the initial evaluation treated resilience primarily as a contingency condition, while adjacent operational and economic pressures ultimately became dominant. The constraint stack evolved, but the architecture remained optimized for the original assumption set.

real even if it does not show up in initial comparisons. Some projects are harmed less by what they spend today than by what they foreclose too early. This creates project optionality, changing the value assessment, which will be addressed specifically in a follow-on piece regarding the “Time Value of Operations.”

This is why the cost of delay and the cost of poor fit are often underweighted. Many organizations are still better at pricing hardware than pricing friction.⁹ Recent market reporting highlights the severity of this gap, noting that unplanned downtime can cost industrial operators millions of dollars per week in lost productivity and idled labor.¹⁰ When these costs are properly modeled, the ‘premium’ for a faster, more flexible solution often disappears.

None of this implies every project requires a premium solution. It does mean that once several constraints are active, decisions should be evaluated against the economics of the real problem, not just equipment cost.




Figure 3. The difference between a poor-fit power decision and a durable one often begins before technologies are compared.

Stop Buying Equipment, Start Solving Constraints

Better distributed power decisions begin before technologies are compared. They begin with problem definition. This does not require a proprietary model but does require discipline, especially under pressure to move quickly or fit an opportunity into an existing internal process.

A more useful screening sequence looks like this:

1. Define the real operational problem: schedule certainty, uptime, flexibility, runtime economics, deployment burden, phased growth, or some combination.
2. Identify which constraints are truly binding. Only some imperfections are decision-shaping.
3. Identify constraint collisions: pinpoint where interacting conditions create a “gating” reality.
4. Identify the value at stake: for example, uptime, schedule, cost control, operational continuity, or all in different proportions.
5. Compare solutions against the real use case rather than default categories. Only after the first four questions does meaningful comparison begin.

This sequence often reveals a subtle truth: the issue is rarely a lack of options. More often, options looked viable only because evaluation criteria were incomplete, compressed, or assumption-sensitive. The analysis eventually reveals complexity that was always there.

⁹ IEA (2025), *Reliability*.

¹⁰ GlobeNewswire (2025), *\$852M Weekly Downtime Costs*.


The Cost of Misclassification

Distributed power decisions often lose quality because projects are categorized too early. Different power problems require different evaluation logic.

- Low-hour backup and low-cost contingency should be evaluated one way.
- Power critical to commissioning, uptime, expansion, or operational viability should be evaluated another.

Neither is inherently better. They are simply different decisions. The challenge arises near category boundaries, where projects initially labeled as “backup,” “temporary power,” or “bridge power” may prove to be critical-path infrastructure problems shaped by timing, flexibility, and execution realities.

Misclassification is costly. A project evaluated as simple insurance may underweight uptime architecture, deployment optionality, runtime consequences, or schedule certainty. A project treated as temporary may make later expansion harder or more expensive. A project treated as “just wait for utility” may absorb months of unpriced schedule drag.



The failure to classify these projects correctly doesn't just lead to budget overruns; it creates a strategic lag that the market is no longer willing to accept. This shift toward owner-controlled, behind-the-meter infrastructure is increasingly becoming the standard response to a congested utility model, as organizations prioritize certainty of power over the traditional, study-heavy interconnection process.¹¹

Complexity Is Not the Problem. Oversimplification Is.

Project teams naturally gravitate toward the cleanest opportunities: simpler projects are easier to explain, approve, and price. Complexity itself is not the enemy. Oversimplifying a project dependent on several interacting realities is.

Frustrating distributed power decisions often result not from impossibility but from early narrowing of the problem, soft assumptions (e.g., optimistic timelines, stable utility queues, predictable loads), or substitutes compared through easier-to-defend criteria than representative ones.

This produces preventable loss of value: buying solutions that do not protect objectives, delaying decisions until options deteriorate, creating avoidable operational burden, or screening out viable paths.

¹¹ Good, S&P Global (2025), *Data Center Developers Turn to Distributed BTM Power*.

A mature evaluation does not ask: “What is the cheapest or most familiar answer?” but instead: “What conditions must this project satisfy to succeed?”

Honesty in framing is more useful than familiarity where timing, reliability, flexibility, site conditions, and organizational realities collide.

Conclusion: A Framework for Long-Term Value Protection

The first failure is often analytical, not mechanical.

Many distributed power projects struggle because the evaluation model is too compressed for the reality it needs to address. Treating distributed power solely as equipment, backup, or CAPEX can filter projects through assumptions that miss the broader enterprise objective. Conversely, treating it as a dynamic problem shaped by timing, site conditions, economics, operations, and execution realities makes the decision clearer – not necessarily easier, but clearer.

Traditional procurement models often simplify to a near-term or semi-static evaluation frame, optimizing for familiar patterns. However, treating long-lived capital assets as static inevitably leads to operational inflexibility, systematically reducing project value as real-world constraints stack against the initial, simplified analysis.

As demonstrated by the collision failures analyzed in this paper, initial cost savings are quickly eroded when a rigid architecture meets dynamic volatility. When operating constraints stack against the project and intersect with an inflexible system, the organization is often left with few viable alternatives. This forces the enterprise into a costly recasting of the asset to handle those constraints, applying downstream pressure on value drivers to either invest more capital retroactively or accept tradeoffs that destroy the economic value of the original analysis.

To mitigate these risks, organizations must replace simplified, siloed modeling with a comprehensive screening path. Evaluating the entire constraint stack up front is a pragmatic approach to risk management that effectively bridges the gap between initial capital deployment and long-term asset management. Furthermore, this approach creates a structured mechanism to value optionality holistically.

Ultimately the initial capital cost of an asset is an incomplete metric if it results in operational inflexibility and pours the future operation in concrete. Investing in architectural flexibility on Day 0 is a deliberate, defensible decision. By implementing a clear, structured screening path, project sponsors move beyond simplistic cost-minimization, equipping themselves with a framework designed to protect long-term value and ensure the capital infrastructure remains an asset over its lifecycle.

That clarity matters. The most consequential decision is rarely which technology is chosen; it is whether the project was framed correctly before any comparison began.

The most expensive mistake isn't the hardware you buy; it's the constraints you ignored.

About Aurelia Technologies

Aurelia Technologies develops distributed power solutions for applications where deployment speed, operational flexibility, site constraints, and evolving power requirements increasingly shape long-term value. The company focuses on industrial systems built for power-constrained, phased, and operationally sensitive environments where infrastructure decisions must be designed around real-world project constraints rather than static, idealized operating conditions.

This paper is part of an ongoing series on distributed power decision-making under real-world constraints.