VERMONT CLEAN POWER BROKERS LLC

The Full Stack™

Behind-the-Meter Prime Power for AI Data Centers

A WHITE PAPER · JULY 2026

The constraint on AI's growth is no longer chips or capital — it is power that can actually be permitted, built, and energized on schedule. Across the industry, projects stall against the same walls: multi-year interconnection queues, grid upgrades that fall on local ratepayers, heavy cooling-water demand, and the community resistance both invite.

The Full Stack is a purpose-built, behind-the-meter power-and-cooling system that lets a data center generate and use its own power on-site — sidestepping the generation-interconnection queue, presenting a flat and benign profile to the local utility, consuming no water in cooling, and engineered to clear local permitting. It is delivered as a standardized ~5 MW block, tiled to match a site of any size, and offered as patent-pending intellectual property, available to acquire outright or to license — deployable by a developer or builder using their own EPC and supply relationships.

Cooling and water, the industry is now solving on its own; deliverable power on schedule, it is not. That is the wall the Full Stack is built to clear.

1. The Power Wall

AI compute demand is compounding faster than the infrastructure that powers it. A year ago the scarce resources were capital and chips; both now flow freely, and the binding constraint has moved to deliverable power. The obstacle is not one wall but several, each on its own multi-year clock, all arriving at once:

Interconnection. Large new loads now wait years in utility queues before they can connect.

Grid upgrades. When the grid must be reinforced to serve a project, those costs increasingly land on local ratepayers — inviting political and regulatory friction.

Water. Conventional liquid cooling can consume large volumes of water; in much of the country that triggers permitting risk, community opposition, and litigation before a project breaks ground.

Volatility. AI training loads swing violently. Pushed onto the grid, those swings stress local transformers and destabilize neighborhoods.

~50%
of U.S. electricity demand *growth* now comes from data centers — the single largest driver.
IEA, 2026

6.6 GW
PJM's shortfall against its own reliability requirements for 2027 — the first in its history.
PJM / Monitoring Analytics, 2026

$130B+
in projects blocked or delayed by community opposition in Q1 2026 alone (75+ projects) — matching all of 2025.
Data Center Watch, 2026

71%
of Americans oppose a data center in their area — more than oppose a nuclear plant.
Gallup, 2026

And efficiency will not relieve the pressure. The Jevons paradox — first observed in coal in 1865 — holds that as a resource becomes cheaper to use, total consumption rises rather than falls. Every gain in compute efficiency has been met with more compute, not less. Demand is not slowing; the walls are simply growing taller.

THE EVALUATION FRAMEWORK HAS ALREADY SHIFTED

Traditional infrastructure was judged on cost: CAPEX, then OPEX, then LCOE. AI infrastructure is judged on speed:

Time-to-Power ›› Time-to-Compute ›› Time-to-Revenue

Reliability protects the revenue. Energy cost finishes last. The logic is one line: **one month of earlier compute is worth more than years of cheaper power.** The Full Stack is engineered to win on speed — permitted, energized capacity on schedule.

2. The Market Has Already Moved

The response taking shape across the most competitive AI campuses is to stop waiting on the grid and build power on-site, behind the meter. What looked contrarian eighteen months ago is now consensus: leading hyperscalers and AI labs are moving to self-generate at campus scale rather than queue for interconnection.

The market has converged on the cooling side as well. In June 2026, NVIDIA published its DSX AI factory reference design for the Rubin generation — the first infrastructure to run 100% liquid cooling, every chip and networking component in a sealed loop with no fans anywhere in the system. The coolant runs at 45 °C, warm enough that outdoor dry coolers reject the heat without mechanical chillers for nearly the entire year in temperate climates. On-site cooling water falls from roughly 2.6 million gallons per megawatt per year to essentially none. Warm-water, closed-loop, zero-consumptive-water thermal design is now the published industry reference.

Capital is abundant and actively hunting for power projects that can actually be permitted and energized. The scarce asset is no longer money — it is a project that clears the walls. The Full Stack is

built to be that project.

3. What the Full Stack Is

The Full Stack is a behind-the-meter prime-power-and-cooling system. One standardized block — approximately 5 MW of critical IT load — holds everything a site needs to run; blocks are tiled together to serve a data center of any size.

> **ON COOLING — AND WHY IT IS NO LONGER THE ARGUMENT**
>
> The Full Stack's warm-water, closed-loop, zero-consumptive-water cooling is the same approach the industry's leading platforms have now adopted as their reference, including NVIDIA's published Rubin DSX design. The cooling is not a bet; it is where the market has already landed. The point is what that frees the conversation to be about: the Full Stack's advantage is not that it cools without water — increasingly, every serious design will — but that it pairs that cooling with on-site prime power that clears the interconnection wall and turns a site on years sooner.

	How the Full Stack answers it
Multi-year wait to connect new power to the grid	**Sidestepped** — the site makes its own prime power and connects as an ordinary utility customer for backup only, skipping the generation-interconnection queue.
Grid upgrades that raise local electric bills	**Avoided** — no new transmission or substation build-out to pass on to ratepayers.
Heavy cooling-water use and permitting risk	**Eliminated** — a sealed loop rejects heat to the air; no towers, no withdrawals, no discharge. This is now the industry-reference approach; the Full Stack meets that standard out of the box.
Compute swings destabilizing the grid	**Contained** behind the meter — the utility sees a flat, quiet load, protecting local transformers.
Community opposition and slow permits	**Designed for social license** — quiet, water-safe, lower-emission, built to clear local review.
Speed of deployment	**Standardized blocks** tiled to scale — repeatable and faster to field than one-off builds.

4. The Modular Advantage

Because the Full Stack is delivered as a standardized block rather than a bespoke design, it is repeatable. A developer with a pipeline of sites can deploy the same engineered block again and again, tiling to each site's load, instead of re-engineering power and cooling from scratch every time. That repeatability compresses schedule, de-risks delivery, and makes capacity planning predictable across a portfolio — the difference between fielding a product and commissioning a series of one-off science projects. For a builder operating at pipeline scale, this is where speed and certainty come from.

It also changes the conversation with the community. A standardized block can be reviewed and approved incrementally, one block at a time, rather than as a single large-block request that a town must accept or reject whole. And because the design repeats, the workforce that builds it can be trained once and deployed again — an economic-development argument a bespoke gigasite structurally cannot make.

5. The Intellectual Property

What makes the Full Stack work is not the components — those are commercially available from established suppliers, and the cooling approach is now industry-standard. The value is the method: how the components are sized and orchestrated to operate as a single, coordinated system — the prime-power, storage, and cooling subsystems dispatched together so the site runs flat, firm, and water-free behind the meter.

That method is the protected, patent-pending intellectual property at the core of the Full Stack, with provisional filings on record and priority dates established. The architecture has been independently reviewed and endorsed by a licensed Professional Engineer — a Cornell-trained mechanical engineer with decades in clean energy. The engineering detail — the how — is disclosed under NDA in a controlled technical exchange; the sizing and dispatch relationships themselves transfer at term sheet.

6. Two Ways to Own It

Vermont Clean Power Brokers is the owner of the Full Stack IP. VCPB does not build, own, or operate facilities. The architecture is offered on two paths, presented on equal footing — each a legitimate, complete transaction, each fitting a different strategy.

Acquire

Own the IP portfolio outright. Control the roadmap and the full market, capture 100% of the value you build on it, and hold it as a clean, one-time transaction.

License

Deploy under license using your own engineering, procurement, and construction resources and your own supply relationships. Vendor-agnostic by design, structured to scale with deployment — block by block, across a pipeline.

Specific terms for either path are set under NDA, matched to the scale and shape of your program. The opportunity is the same either way; only the structure differs.

7. Bull or Bear — The Thesis Holds Either Way

There is serious debate, bordering on argument, that the AI buildout is overbuilt — that the models are inefficient, that capital has been substituted for engineering discipline, and that a correction is coming. We don't need you to reject that argument. The Full Stack is built to survive it.

If the buildout accelerates, the Full Stack sells speed into a land rush. Time-to-power is the binding constraint, capacity carries a premium, and the architecture that skips the interconnection queue wins on schedule.

If the buildout corrects, capital gets disciplined. It stops funding speculative gigasites and starts asking the only question that survives a correction: can this actually be permitted, energized, and delivered? That is precisely the question the Full Stack is engineered to answer. The projects that die in a correction are the ones that could never clear the walls. The ones that survive are the ones that can be built.

The layer beneath doesn't change. Models will get more efficient — they should. Chips will change. Architectures will be replaced. The advantage the Full Stack holds and will retain is that every one of them still requires electrons delivered, heat rejected, and a permit signed. In short: the Full Stack powers any model, and **the power does not care about the algorithm.**

And the precedent isn't hypothetical. Every efficiency gain in the history of computing has produced more computing, not less demand for power. An adjacent market makes the same point: when solar modules became dramatically more efficient, the racking and siting businesses didn't shrink — the market grew faster than the efficiency gain. The Full Stack is more than just the smart play — it's the longevity play.

8. Why Now

The window is open and finite. The walls are real, the market has accepted that behind-the-meter power is the answer, and capital is waiting for projects that can be permitted and energized. Cooling and water are now being solved across the industry; deliverable power on schedule is not — and that is the constraint that decides who gets compute online first. The Full Stack exists today, as protected IP, ready to acquire or license and deploy. The advantage belongs to those who move while the field is still forming.

> **Stated with discipline.** The Full Stack is offered as specified, patent-pending intellectual property. It is not represented as an operating facility, and no deployments, pilots, or revenue are claimed. The asset is the architecture, the specification, and the filed IP — complete, defensible, and ready. We would rather understate and be verified than overstate and be corrected.

Bob Lewis

Founder & CEO · Vermont Clean Power Brokers LLC

802-230-6561 · learnmore@VTCleanPower.com · vtcleanpower.com/fullstack

Patent Pending — U.S. Provisional Ser. 64/029,433, filed April 4, 2026. Market figures per cited public sources, 2026, and verified July 2026. NVIDIA, Rubin, and DSX are trademarks of NVIDIA Corporation, referenced here as published industry reference material; no affiliation or endorsement is claimed or implied. The Full Stack™ is offered as intellectual property; no operating facility, deployment, or revenue is claimed. Engineering detail and the protected method are disclosed only under executed NDA.