Flux Core Data Systems and Energy

AI Infrastructure White Paper Series

July 2026

Own the Energy, Monetize the Data.

Table of Contents

1. Offtake Agreements: The Contract That De-Risks a Megawatt

Keywords: bankable megawatts, contracted AI demand, offtake financing, capacity reservation, anchor tenant compute

Executive Summary

A megawatt of planned data center capacity is not an asset until someone has agreed to pay for the power it will deliver. That distinction has become the central question in AI infrastructure finance, because the supply of new generation is falling far behind the demand that AI is creating. Bank of America estimates that the United States needs more than 230 GW of new generating capacity over the next five years, while regulated utilities plan to add only about 93 GW of accredited supply, leaving a gap of more than 100 GW, and data centers alone account for roughly 125 GW of new load (Source 1). In a market that short of supply, the scarce and financeable thing is a contract for delivered power. Flux Core Data Systems and Energy leads every deployment with a signed offtake position, currently a Memorandum of Understanding covering 50 to 100 MW of demand, and continues to open capacity to neoclouds, enterprises, universities, hospitals, municipalities, national labs, and AI platforms. This paper explains why the offtake agreement is the instrument that de-risks a megawatt, and how Flux Core makes contracted capacity deliverable with power agnostic generation and closed loop cooling that uses zero outside water.

The platform behind that promise is the Nexus R1000, a 1MW containerized data center that supports powered shell and bare metal loads up to 1.3MW of IT capacity, from H200 and B200 class racks at roughly 40 to 60 kW to GB300 NVL72 class racks at roughly 130 to 140 kW. Every unit is built to Tier III design standards targeting 99.982 percent uptime, is relocatable, and comes online in three to six months. That combination is what allows a signed offtake to translate into real, delivered megawatts on a schedule a lender can trust, rather than a promise contingent on a grid connection that may never arrive on time.

The Problem

The supply gap is now measurable and large. Bank of America projects a demand compound annual growth rate of 4.1 percent from 2026 to 2030 and warns that the roughly 93 GW of accredited capacity utilities plan to add cannot cover the more than 230 GW the system will need, a shortfall of over 100 GW (Source 1). A shortfall of that size changes the economics of every project. When capacity is abundant, a data center is underwritten as real estate, valued on the expectation that tenants will eventually fill it. When capacity is scarce, the same building is underwritten on a single question: is the power contracted and deliverable?

That question is sharper because the constraint is not only how much generation gets built. It is also how fast load can connect and whether the grid can carry it. The PJM interconnection, the largest grid market in the country, now reports that data centers account for 94 percent of peak load growth through 2030 (Source 2). When one category of customer drives nearly all incremental demand, ordinary buyers can no longer assume they will find capacity waiting for them. They must reserve it in advance or go without.

For a lender or an equity investor, a speculative build carries the entire risk that the developer never signs paying customers, that the interconnection slips, and that expensive equipment sits idle while the demand it was ordered for goes elsewhere. Capital either prices that risk punitively or declines the project outright. The offtake agreement resolves the problem by substituting a contractual revenue stream for a market bet. It converts a speculative build into a financeable one, which is why contracted demand has become the true unit of value.

The mismatch runs in both directions. Developers cannot finance without committed demand, and buyers cannot secure guaranteed capacity because everyone is competing for the same scarce, deliverable megawatts. Capacity without committed demand is unfinanceable, and demand without deliverable capacity is unfulfillable. A single instrument, the offtake agreement, is what closes both gaps at the same time, provided the capacity behind it can actually be delivered on a timeline the buyer can use.

It is worth being precise about what the Bank of America numbers imply for timing. A demand growth rate of 4.1 percent compounding from 2026 through 2030, layered on top of a base that already strains supply, means the gap does not close on its own and in fact widens in the near term (Source 1). Utilities plan in decade long increments and build accredited capacity slowly, while AI demand arrives in quarters. That structural difference in tempo is why the market cannot simply wait for balance to return. Every quarter that passes with 125 GW of data center load pressing against 93 GW of planned utility additions makes contracted, deliverable capacity more valuable, not less, and it rewards the buyers and developers who moved early with a signed agreement in hand.

Real-World Example

The public data makes the stakes concrete. With utilities planning to add only about 93 GW against a need above 230 GW (Source 1), and data centers driving 94 percent of PJM peak load growth (Source 2), a buyer that waits for open market capacity to appear is likely to be disappointed. The parties that win are the ones that lock capacity ahead of the shortfall.

Consider an illustrative scenario. A regional neocloud has signed inference customers and a growing training pipeline but no guaranteed way to power the GPUs its customers expect. Going to the public grid means entering a queue behind the data center load that already dominates connection growth, an outcome that would force the neocloud to turn away revenue today. Instead it reserves a defined block of megawatts inside a Flux Core deployment, with delivery measured in months rather than years. This scenario is illustrative and does not describe a specific Flux Core customer, yet it mirrors the exact behavior the supply gap predicts. That reservation becomes the demand signal that anchors the site and makes it financeable.

How Flux Core Solves It

Flux Core makes megawatts bankable by making them diversified and deliverable. The demand book is deliberately spread across neoclouds, enterprises, universities, national labs, hospitals, municipalities, and AI platforms, so no single sector dictates the timeline or the credit risk. A diversified offtake book behaves like a diversified revenue portfolio and earns better terms than a single tenant bet. When one sector softens, the others carry the deployment, which is the resilience capital markets reward.

Deliverability rests on engineering choices that address the exact constraints the market data identifies. First, the Nexus R1000 is a 1MW containerized data center that comes online in three to six months rather than the multi year timelines of conventional construction. Reserved capacity that arrives in months is a different product from capacity promised for years out. Second, Flux Core is power agnostic. Reserved capacity can sit on stranded or associated natural gas, on renewable solar paired with battery energy storage, on a hybrid microgrid, or on an existing grid tie. A commitment therefore never stalls behind the interconnection crowding that data center load has created in markets like PJM (Source 2).

Underpinning it all is a closed loop liquid cooling system that uses zero outside water. Because the fluid recirculates with no cooling tower and no evaporative loss, there is no make up water permit to secure and no dependence on local water supply. That removes an entire category of permitting and siting risk from the delivery path, which is exactly what a lender or an anchor tenant needs to see before signing. The Nexus platform targets Tier III design standards and 99.982 percent uptime, so contracted capacity is not only deliverable but dependable. In a market short more than 100 GW of accredited supply (Source 1), that combination turns a contracted megawatt into collateral.

The result is a reframing of what an offtake partner is actually buying. It is not square footage and a hope that power eventually arrives. It is a defined block of capacity with a defined delivery date, an energy path that does not depend on the interconnection crowding that data center load has created (Source 2), and a cooling architecture that removes the water and permitting risks that derail conventional projects. For a lender, that combination is what converts a contracted megawatt into collateral it can underwrite. For a buyer, it is what converts a reservation into a guarantee. The Memorandum of Understanding covering 50 to 100 MW is the proof point that serious counterparties already treat Flux Core capacity as bankable, and the remaining capacity is open to partners who want the same certainty in a market that is only growing tighter.

None of this asks a partner to change its business model. A neocloud keeps reselling GPU hours, an enterprise keeps training its private models, a hospital or a national lab keeps its mission, and a university keeps its research agenda. What changes is that the megawatts underneath all of them are contracted, deliverable, and cooled without a drop of outside water, sitting on a power path that does not wait in a queue crowded by hyperscale load (Source 2). That is the quiet advantage of an offtake led model. It de-risks the megawatt for the developer and the lender at the same time it de-risks supply for the buyer, which is why Flux Core treats the signed agreement, not the building permit, as the first step of every project.

Conclusion and Call to Action

When the United States needs more than 230 GW and utilities will add only about 93 GW (Source 1), and when data centers drive 94 percent of the largest grid market's peak growth (Source 2), the offtake agreement is the instrument that turns scarce megawatts into bankable demand. Flux Core builds every deployment around that instrument, backed by containerized delivery in months, a power agnostic energy path, and cooling that draws zero outside water. If you need guaranteed AI capacity, or want to anchor a site as an offtake partner, contact Flux Core to see which megawatts remain open. Own the Energy, Monetize the Data.

Sources

1. Utility Dive, Bank of America on the United States generation gap and data center load.
https://www.utilitydive.com/news/ai-data-center-growth-utilities-generation-plans/825541/

2. Matter Fact, PJM data centers driving 94 percent of peak load growth through 2030.
https://www.matterfact.com/newsletter/2026-07-07-data-centers-pjm-94-percent

2. Bare Metal Provisioning: Engineering the Shell for 130 kW Racks

Keywords: high density powered shell, 1.3MW IT load, 130 kW racks, direct liquid cooling, GPU ready bare metal

Executive Summary

The economics of an AI data center are now decided at the rack, not at the fence line. Cooling has become the fastest growing line in the power budget, and the racks themselves draw loads that would have described an entire row a decade ago. Gartner forecasts that cooling electricity will rise 22.6 percent to 195 TWh in 2026, part of a jump in global data center electricity from 447 TWh in 2025 to 565 TWh in 2026, a 26 percent increase, with AI optimized servers nearly doubling from 95 to 175 TWh (Source 1). A powered shell designed for a few kilowatts per cabinet cannot absorb that thermal reality. Flux Core Data Systems and Energy has re-engineered the shell around it, delivering a bare metal environment provisioned for up to 1.3MW of IT load, ready for GB300 NVL72 class racks at roughly 130 to 140 kW, cooled by a closed loop liquid system that consumes zero outside water, powered by a power agnostic energy path, and deployable in three to six months. This paper explains why the thermal math breaks conventional shells and how Flux Core solves density at the foundation.

The foundation is a purpose built platform rather than a retrofitted room. The Nexus R1000 arrives as a containerized 1MW system provisioned for up to 1.3MW of IT load, with liquid cooling plumbed to the chip and a power path selected to match the site. Because the unit is relocatable and comes online in three to six months, a tenant can energize dense racks now and keep the option to move capacity as its hardware roadmap and customer base evolve. Density, cooling, power, and speed are treated as a single engineered system, which is the only way to serve racks that draw as much as 80 to 100 homes without rebuilding a hall from scratch.

The Problem

Rack density has crossed a threshold that air cooling cannot follow. A single high density AI rack can now draw as much electricity as roughly 80 to 100 homes, and Gartner attributes a 22.6 percent rise in cooling electricity, to 195 TWh in 2026, largely to this shift toward accelerated computing (Source 1). Air cooling can reasonably serve racks in the range of ten to twenty kilowatts. Beyond that, moving enough air to remove the heat becomes physically impractical and wasteful. Modern accelerators blow past that ceiling immediately.

An H200 or B200 class deployment runs at roughly 40 to 60 kW per rack, and a GB300 NVL72 class rack reaches 130 to 140 kW.

A conventional powered shell was designed and permitted around evaporative cooling towers and air handling for a much lighter load profile. It cannot be retrofitted to those densities without effectively rebuilding it. That is the trap. Even where the shell exists and the power is available, the hall cannot cool the racks the tenant needs to install. The cooling ceiling, not the floor space, is the binding constraint.

The power side compounds the problem. Gartner projects global data center electricity rising 26 percent in a single year and AI optimized servers reaching 31 percent of data center power in 2026 (Source 1). A shell that depends on a grid connection inherits the interconnection crowding that AI load has created, so the tenant faces two failures at once: a hall that cannot cool dense racks and a power path that may take years to energize. Solving one without the other accomplishes nothing.

The result is a category of stranded demand. Tenants have hardware, capital, and signed customers, but no dense, powered, deploy ready white space to put the hardware into. High end accelerators depreciate quickly and the contracts that justify them have hard start dates, so a tenant that cannot energize dense racks on time is paying for idle silicon while competitors serve the market. The conventional colocation industry, built for a slower and less dense era, is structurally unprepared to close that gap.

The cooling gap deserves particular emphasis because it is the hidden reason conventional shells fail, and the Gartner data quantifies how quickly it is worsening. Cooling electricity rising 22.6 percent to 195 TWh in a single year, while AI optimized servers climb from 95 to 175 TWh and reach 31 percent of all data center power (Source 1), describes a thermal load that is growing far faster than the halls built to reject it. A powered shell designed and permitted around evaporative cooling towers cannot absorb that curve by adding fans or chillers at the margin. The physics of air cooling caps out well below the density that a GB300 NVL72 class rack requires, so the only real answer is to change the heat rejection method itself, from air moving through the room to liquid moving to the chip. A shell that was not designed for liquid from the start is, in practice, a shell that has to be rebuilt.

Real-World Example

The public numbers frame the gap precisely. With cooling electricity up 22.6 percent to 195 TWh and AI servers nearly doubling to 175 TWh in a single year (Source 1), a tenant that needs dense racks fast has almost no conventional option. The cooling and power provisioning of an existing shell were designed for a load profile that no longer applies.

Consider an illustrative scenario. A GPU tenant has secured an allocation of GB300 NVL72 class systems and signed customers who expect capacity this year. Evaluating a conventional powered shell, the tenant discovers two disqualifying facts: the hall cannot cool racks anywhere near the 130 to 140 kW the hardware requires, and the power path behind the shell is years from energization. This scenario is illustrative and does not describe a specific Flux Core customer. It reflects the structural mismatch the data describes, and it is exactly the situation the Flux Core powered shell is built to resolve.

How Flux Core Solves It

Flux Core delivers a bare metal environment provisioned for up to 1.3MW of IT load, pre engineered for the densest AI and GPU workloads. Density is a cooling problem before it is a power problem, and Flux Core solves cooling at the foundation. The closed loop liquid cooling system recirculates a single fluid and consumes zero outside water, which eliminates cooling towers, make up water permits, and the evaporative losses that cap traditional halls. That thermal headroom is what makes extreme density possible while the industry cooling bill climbs 22.6 percent (Source 1). The platform supports Configuration 2 for H200 and B200 class systems at roughly 40 to 60 kW per rack, and Configuration 3 for GB300 NVL72 class systems at roughly 130 to 140 kW per rack, all at Tier III design reliability targeting 99.982 percent uptime.

The shell is also power agnostic, which breaks the timeline problem. Whether a site offers stranded or associated natural gas, renewable solar with battery energy storage, a hybrid microgrid, or an existing grid tie, Flux Core builds the matching power path rather than waiting in a queue crowded by the very AI load that Gartner shows growing 26 percent in a year (Source 1). That is how a 1.3MW environment reaches edge and constrained sites conventional colocation avoids, and why deployment happens in three to six months.

The container form factor ties it together. Because the Nexus R1000 is containerized and relocatable, a tenant is not committing capital to a fixed building on a fifteen year horizon against a market that changes every quarter. Tenants drop in their own hardware or lease

Flux Core hardware, energize dense racks in months, and retain the option to relocate as needs shift. Cooling is engineered in, power certainty is engineered in, and speed is engineered in, so a tenant no longer has to choose between fast, dense, and powered.

It is worth stating plainly what the rebuilt shell replaces. The old model asked a tenant to accept a slow, capital heavy, grid dependent building that could not cool modern hardware, and to hope that power and cooling could be sorted out later. The Flux Core model makes cooling and power provisioning first class design elements rather than afterthoughts. Liquid cooling engineers density in, the power agnostic architecture engineers power certainty in, and the containerized form factor engineers speed in. As the industry cooling bill climbs 22.6 percent and AI server power nearly doubles in a year (Source 1), that difference is not cosmetic. It is the difference between a hall that can host a 130 kW rack and one that physically cannot, delivered on a timeline measured in months rather than the years a grid dependent rebuild would demand.

The economics follow directly from the engineering. A tenant that can energize a 130 kW rack in months rather than years puts expensive accelerators to work before they depreciate and serves customers on the start dates their contracts specify. Idle silicon is the single largest avoidable cost in an AI deployment, and it is created entirely by the gap between when hardware arrives and when a hall can actually cool and power it. By closing that gap with liquid cooling, a power agnostic energy path, and a containerized form factor, Flux Core converts the density that breaks conventional shells into the density that a tenant can monetize, which is the whole point of provisioning bare metal for AI in the first place.

Conclusion and Call to Action

When cooling electricity climbs 22.6 percent to 195 TWh and AI optimized servers nearly double in a single year (Source 1), the shell has to be engineered around 130 kW racks rather than retrofitted for them. Flux Core delivers that shell: up to 1.3MW of IT load, GB300 NVL72 class density, closed loop cooling that uses zero outside water, and a power agnostic energy path that energizes in months. If you are provisioning dense GPU racks on a deadline, ask Flux Core about a liquid cooled powered shell ready for up to 1.3MW. Own the Energy, Monetize the Data.

Sources

1. Tom’s Hardware, Gartner forecast on AI server power, data center electricity, and cooling growth. https://www.tomshardware.com/tech-industry/artificial-intelligence/ai-servers-will-consume-more-power-than-conventional-data-center-hardware-by-2027-gartner-forecasts

2. Matter Fact, behind the meter power and data center demand context. https://www.matterfact.com/newsletter/2026-07-03-data-centers-done-waiting-grid

3. Power as a Service: Paying for Speed Without Paying the Premium

Keywords: power as a service, speed to power, on-site generation, behind the meter power, PaaS microgrid

Executive Summary

The fastest way to power an AI cluster is also, right now, the most expensive. On site and behind the meter power is running roughly an 80 percent premium to the grid, about 140 to 150 dollars per MWh against near 80 dollars, and the heavy duty turbines that would relieve the shortage are sold out at the big three manufacturers (Source 1). Speed to power has a price, and most developers are paying it because they have no alternative. Flux Core Data Systems and Energy offers a different arrangement. Under Power as a Service, Flux Core finances, owns, and operates the generation, and the customer buys reliable power for compute as a clean operating expense. The model is now going mainstream at scale: Brookfield and Bloom Energy expanded a fuel cell financing partnership from 5 billion to 25 billion dollars, covering up to 1 GW of on site fuel cells across an AI data center portfolio (Source 2). This paper explains the speed to power premium and how Flux Core delivers speed without forcing the customer to carry that premium or the capital behind it.

The delivery vehicle is the Nexus R1000, a 1MW containerized data center that pairs the generation Flux Core owns and operates with liquid cooling that uses zero outside water. Because the platform is power agnostic and relocatable, Flux Core can match the least expensive reliable energy resource at each site, whether that is stranded gas, solar with battery storage, a hybrid microgrid, or a grid tie, and stand the system up in three to six months. The customer sees only a predictable rate for delivered, Tier III grade power, while the complexity of sourcing, financing, and operating that power stays with Flux Core.

The Problem

The grid can no longer connect AI load fast enough, and the market has repriced speed accordingly. On site and behind the meter power now commands roughly an 80 percent premium over grid power, about 140 to 150 dollars per MWh compared with near 80 dollars per MWh, and the largest turbine manufacturers are sold out of heavy duty units, extending lead times further (Source 1). A developer that needs power this year faces a choice between a multi year grid queue and a premium priced on site option, with little room in between.

That premium exists because the alternative is worse. Interconnection timelines have stretched as data center load has come to dominate grid growth, and the equipment needed to add generation is backlogged. When the only fast path to power is scarce, its price rises. Developers pay because idle GPUs cost more than expensive electricity, but the premium erodes the economics of the whole project and forces capital into generation assets the developer never wanted to own.

There is also a capital structure problem. Building and owning generation, switchgear, and storage locks a compute operator into an energy business it has no expertise in and no desire to run. The capital sits on the balance sheet, the operations distract from the core product, and the timeline is still hostage to equipment lead times. The developer ends up paying both the speed premium and the cost of becoming an accidental power company.

The market is voting with its money on a different answer. When Brookfield and Bloom Energy expand an on site power financing partnership fivefold, from 5 billion to 25 billion dollars, to deploy up to 1 GW of fuel cells across an AI portfolio (Source 2), it signals that on site generation is becoming the default rather than the exception. The open question for a compute operator is not whether to go behind the meter. It is how to do so without carrying the premium and the capital.

The premium is also a moving target that tends to rise with demand. As long as heavy duty turbines remain sold out at the big three manufacturers (Source 1), the scarcity that props up the 80 percent premium persists, and a developer negotiating today has little leverage over the price. Worse, a developer that commits capital to buying and operating its own generation at these prices locks in a cost structure just as the underlying equipment market is at its tightest. The compute operator ends up exposed twice, once to the premium on the power itself and again to the risk that it has overpaid for generation assets whose value depends on a scarcity that eventually eases. What the operator actually wants is delivered power at a predictable price, insulated from both the premium and the equipment market, which is precisely what an ownership model cannot provide and a service model can.

Real-World Example

The pricing data makes the stakes clear. At an 80 percent premium, roughly 140 to 150 dollars per MWh against near 80 dollars for grid power, and with turbines sold out (Source 1), a developer buying on site power at market rates pays a heavy toll for speed. Meanwhile the scale of the Brookfield and Bloom commitment, up to 1 GW and 25 billion dollars of financing (Source 2), shows how much capital is now chasing on site generation.

Consider an illustrative scenario. An enterprise needs a dedicated inference cluster online within the year and cannot wait for a grid connection. Buying on site power at the market premium would raise its cost per token and force it to finance generation equipment it does not want. Instead it contracts with Flux Core for Power as a Service, paying a predictable rate for delivered power while Flux Core owns and operates the underlying generation. This scenario is illustrative and does not describe a specific Flux Core customer, but it reflects the exact trade the market data describes.

How Flux Core Solves It

Under Power as a Service, Flux Core finances, owns, and operates the generation and distribution, and the customer buys reliable power for compute as a clean operating expense. There is no capital locked into generators, switchgear, or storage, and no project timeline held hostage by a utility queue. The customer avoids becoming an accidental power company while still getting the speed that on site generation provides.

Because Flux Core is power agnostic, the system under the customer's load can be stranded or associated natural gas, renewable solar paired with battery energy storage, a hybrid microgrid, or a grid connection, sized to the customer's ramp and location. That flexibility is what lets Flux Core source the least expensive reliable power available at a given site rather than defaulting to the premium priced options the broader market is paying (Source 1). It also aligns Flux Core with the on site generation wave that Brookfield and Bloom are financing at scale (Source 2), while structuring the arrangement so the customer never touches the capital.

Thermal efficiency holds up alongside it. Closed loop liquid cooling with zero outside water keeps efficiency high and takes the facility off local water supply, a decisive advantage in the arid, low cost power regions where compute increasingly wants to live. The Nexus R1000 delivers this as a 1MW containerized data center, built to Tier III design standards, deployable in three to six months. The customer pays a predictable rate while Flux Core monetizes the underlying energy structures, so speed to power stops being a premium the customer pays and becomes a service the customer buys.

The structural advantage of the service model becomes clearest when the two paths are compared side by side. Under the ownership path, a compute operator raises capital, buys generation equipment into a sold out market, waits on turbine lead times, staffs an energy operations team, and still bears the full premium on the power (Source 1). Under Power as a Service, the operator signs a contract, receives delivered power at a predictable rate, and keeps its capital and its attention on the compute product that actually earns revenue. Flux Core absorbs the equipment risk, the operations burden, and the fuel flexibility decisions, spreading them across a portfolio in the same way the Brookfield and Bloom partnership spreads on site generation across an AI fleet (Source 2). The customer captures the speed of behind the meter power without inheriting the balance sheet or the premium that usually comes with it.

The arrangement also changes the risk profile of the whole project. When power is an operating expense at a predictable rate, a compute operator can price its own product with confidence, forecast its margins, and avoid the balance sheet exposure of owning generation into a sold out equipment market (Source 1). Flux Core carries the fuel flexibility, the equipment procurement, and the operations, and spreads those risks across a portfolio rather than concentrating them in a single site. The customer gets the one thing it actually needs, reliable power delivered where and when the compute runs, without becoming the accidental power company that the ownership path forces it to be.

Conclusion and Call to Action

On site power at an 80 percent premium with turbines sold out (Source 1) is the price of speed in today's market, and the Brookfield and Bloom expansion to 25 billion dollars shows how mainstream behind the meter generation has become (Source 2). Flux Core lets a customer capture that speed without carrying the premium or the capital, through a power agnostic, zero water platform delivered as an operating expense. If you have been told to bring your own power, ask Flux Core how Power as a Service delivers it. Own the Energy, Monetize the Data.

Sources

1. Matter Fact, speed to power premium, on site pricing, and turbine backlog.
https://www.matterfact.com/newsletter/2026-07-03-data-centers-done-waiting-grid

2. Construction AI Brief, Brookfield and Bloom Energy expand on site fuel cell financing to 25 billion dollars. https://constructionaibrief.com/posts/2026-07-02-brookfield-bloom-national-grid-onsite-power-turbine-backlog

4. Oil and Gas Well Deployments: The Molecules the Pipeline Left Behind

Keywords: stranded gas to compute, flare gas monetization, wellhead data center, associated gas power, remote well AI

Executive Summary

A pipeline can rescue a large flare, but it cannot reach every well. That is the nuance in the latest global flaring data, and it defines a real opportunity. The World Bank 2026 Global Gas Flaring Tracker reports that flaring rose for a third straight year to 167 billion cubic meters in 2025, the highest level since 2019, wasting an estimated 54 billion dollars and releasing roughly 429 million tons of carbon dioxide equivalent (Source 1). Yet within that global rise, the United States actually cut flaring 7 percent to about 5 billion cubic meters as new Permian pipeline capacity came online (Source 2). The lesson is precise. Pipelines helped the big, accessible flares, but takeaway is scarce and slow to build, so stranded gas persists at smaller and more remote wells. Flux Core Data Systems and Energy deploys containerized data centers at those wells, converting stranded and associated gas into compute revenue on site, cooled by a closed loop liquid system that uses zero outside water. This paper explains why the pipeline leaves molecules behind and how Flux Core puts them to work.

The unit that does the work is the Nexus R1000, a 1MW containerized data center engineered to run at the edge of infrastructure. It trucks to a remote pad, uses wellhead gas for on site generation, rejects heat through a closed loop that needs no outside water, and relocates when the field declines, all while operating to Tier III design standards. Because it is power agnostic, the same unit can transition from wellhead gas to a renewable or grid source later without being stranded itself. That mobility and fuel flexibility are what let compute finally reach the small and remote wells that a fixed pipeline was never going to serve.

The Problem

Flaring is not a marginal problem. The World Bank reports 167 billion cubic meters flared globally in 2025, the third consecutive annual increase and the highest since 2019, wasting about 54 billion dollars in energy and emitting roughly 429 million tons of carbon dioxide equivalent, including about 50 million tons from unburned methane (Source 1). Nine countries account for more than 80 percent of the total. The gas is burned because there is no economic way to move it from where it is produced to where it can be used.

The United States figure adds the crucial detail. American flaring dropped 7 percent to roughly 5 billion cubic meters in 2025, and the reason was new Permian pipeline capacity that finally gave large producers a path to market (Source 2). That is good news, but it is also a narrow lesson. Pipelines are capital intensive and slow to permit and build. They get built to serve the largest, most concentrated volumes first. The smaller wells, the older wells, and the remote wells far from a trunk line remain stranded, because the economics of a dedicated pipeline never close for them.

So the molecules the pipeline leaves behind keep burning. An operator at a remote pad watches associated gas go up the flare stack because no takeaway exists and none is coming, while the same molecule could be running high margin compute. The gas is a liability to be disposed of rather than an asset to be sold, and every year of flaring is a year of wasted energy and avoidable emissions at exactly the wells the pipeline solution does not reach.

Water compounds the siting challenge in producing basins, which are frequently arid. A conventional data center depends on large volumes of water for evaporative cooling, which is precisely what a remote well pad cannot supply. Any solution that brings compute to the wellhead must therefore solve cooling without outside water, or it simply trades one scarcity for another.

The scale of what is left behind is easy to underestimate when the headline number for one country improves. The United States decline of 7 percent to about 5 billion cubic meters is real, but it sits inside a global total of 167 billion cubic meters that rose for a third consecutive year (Source 1, Source 2). Even within the United States, the pipeline driven improvement is concentrated in the largest, most accessible fields. A pipeline is an expensive, long lived asset that only gets built where volumes are large and durable enough to justify the capital, so the long tail of smaller and declining wells, and the remote wells far from any trunk line, is systematically excluded from that solution. Those wells do not stop producing associated gas. They simply keep flaring it, which is why on site conversion, rather than midstream takeaway, is the only economic path for the molecules the pipeline was never going to reach.

Real-World Example

The data draws the map. Global flaring at 167 billion cubic meters and rising for a third year (Source 1) shows how much gas has no path to market, while the United States 7 percent drop tied to new Permian pipeline capacity (Source 2) shows that pipelines help only where they reach. Between those two facts sits a large population of smaller and remote wells that will keep flaring because a pipeline will never be built to serve them.

Consider an illustrative scenario. An operator holds acreage with several smaller wells producing associated gas that no midstream company will take, so the gas is flared under permit. This scenario is illustrative and does not describe a specific Flux Core customer. The operator brings a containerized Flux Core unit to the pad, uses the wellhead gas to generate power on site, and converts what was a flared liability into compute revenue. The unit relocates as the field’s production curve shifts, so the operator is not building fixed infrastructure against a declining well.

How Flux Core Solves It

Flux Core deploys containerized data centers at the wellhead and converts stranded and associated gas into compute revenue on site, with no midstream takeaway required. Gas that would be flared, vented, or sold at a steep discount becomes fuel for on site generation that powers high margin compute and token output. That is a new revenue line off an existing asset, and it directly targets the molecules the pipeline left behind (Source 1, Source 2).

The platform is power agnostic, which is what makes the wellhead deployment durable. The unit runs on wellhead gas today, with a clear path to renewables or a grid tie later as the field matures or infrastructure arrives. The operator is never locked into a single fuel source, and the deployment adapts as the well’s output changes. Because the Nexus R1000 is containerized and relocatable, it trucks to a remote pad, energizes in months, and moves to the next opportunity when the current one winds down.

Closed loop liquid cooling with zero outside water is what makes any of this possible in an arid basin. The system recirculates its fluid with no cooling tower and no evaporative loss, so a remote pad with no water supply is no obstacle. The deployment brings its own cooling as well as its own fuel, which is why it can reach wells that conventional data centers never could. Built to Tier III design standards targeting 99.982 percent uptime, the unit turns a stranded liability into dependable, revenue generating compute while shrinking the flare and improving the emissions profile of the site.

The environmental case reinforces the economic one rather than competing with it. Flaring wastes an estimated 54 billion dollars a year and releases roughly 429 million tons of carbon dioxide equivalent, including about 50 million tons from unburned methane that escapes when flares operate inefficiently (Source 1). Using that gas in a controlled generator to power compute burns it more completely than an open flare and puts the energy to productive use, so the same molecule that was a pure emissions liability becomes both a revenue source and a smaller environmental footprint. For an operator facing tightening flaring regulations and public scrutiny, a deployment that reduces the visible flare while creating a new revenue line is a rare alignment of the commercial and the regulatory incentive, and it is available today at exactly the wells the pipeline solution passed over.

For the operator, the appeal is that the deployment fits the natural life cycle of a well rather than fighting it. A wellhead produces the most associated gas early and tapers over time, and a relocatable unit can follow that curve, arriving when the gas is abundant and moving to the next pad when it is not. There is no stranded capital in a fixed building tied to a declining well, and no dependence on a pipeline that may never be built. The molecules that the pipeline left behind become a mobile, revenue generating asset, cooled with zero outside water and powered by the gas itself, which turns a disposal problem into a durable line of business.

Conclusion and Call to Action

Global flaring rising to 167 billion cubic meters (Source 1) while United States pipelines cut only the largest Permian flares (Source 2) proves that stranded gas persists exactly where takeaway does not reach. Flux Core turns those left behind molecules into compute at the wellhead, power agnostic and cooled with zero outside water, on a containerized platform that relocates with the field. If you are flaring gas the pipeline will not take, Flux Core can turn it into compute. Let us scope a wellhead deployment on your acreage. Own the Energy, Monetize the Data.

Sources

1. World Bank, 2026 Global Gas Flaring Tracker, flaring rises for a third year.
https://www.worldbank.org/en/news/press-release/2026/06/23/global-gas-flaring-rises-energy-security

2. Payne Institute at Colorado School of Mines, analysis of the 2026 flaring tracker including the United States decline. https://payneinstitute.mines.edu/the-world-bank-2026-global-gas-flaring-tracker-report-shows-a-rise-in-flaring-for-the-third-year-in-a-row/

5. Zero Carbon Footprint: The Water Line Nobody Reports

Keywords: water free liquid cooling, low water AI data center, closed loop cooling, renewable powered compute, resource efficient PUE

Executive Summary

Most sustainability reports for AI infrastructure lead with carbon and quietly omit water, yet water is where the resource story is most visible and most fraught. A single large AI data center can consume up to 5 million gallons of water a day, United States data centers directly consumed 17.4 billion gallons in 2023, and indirect water use through power generation ran roughly twelve times higher, about 211 billion gallons (Source 1). The pressure peaks precisely when communities can least afford it, during the heatwaves when both cooling demand and local water scarcity spike (Source 2). Flux Core Data Systems and Energy treats water as the line that must be engineered to zero, not offset on paper. Our closed loop liquid cooling recirculates its fluid and consumes zero outside water, an approach the industry is now validating: NVIDIA introduced a closed loop liquid cooling design that pushes on site water use toward zero (Source 3). This paper explains why water is the honest measure of a data center's footprint and how Flux Core engineers both water and carbon down rather than buying credits after the fact.

The engineering starts with the Nexus R1000, a 1MW containerized platform whose closed loop liquid cooling recirculates a single fluid with no cooling tower and no evaporative loss, so the direct water draw is zero. Paired with a power agnostic energy path that can prioritize renewables and flare gas over grid power, the same unit shrinks the far larger indirect water footprint embedded in generation. Built to Tier III design standards and deployable in three to six months, it delivers a footprint that is engineered into the hardware and the energy choice rather than assembled from credits after the fact.

The Problem

Water is the resource cost that rarely appears in the headline. Lawrence Berkeley National Laboratory data shows United States data centers directly consumed 17.4 billion gallons of water in 2023, while the indirect water embedded in the electricity they used was roughly twelve times larger, about 211 billion gallons (Source 1). A single large AI facility can draw up to 5 million gallons a day, the equivalent of a small town, largely to feed evaporative cooling towers that boil water away to reject heat.

The timing makes it worse. Water demand from data centers peaks during heatwaves, exactly when rivers run low and communities face restrictions (Source 2). A facility that competes with a town for drinking water during a drought is a political and regulatory liability, not just an environmental one. As AI load grows, the water it consumes, directly and indirectly, becomes a visible flashpoint in the communities that host it.

Carbon accounting has its own honesty problem. A data center can purchase renewable energy credits and still draw fossil power from the grid around the clock, claiming to be clean on paper while burning fossil fuel in practice. Offsets and unbundled credits let an operator report a low carbon number without changing the physical energy flowing into the building. The result is a footprint that looks better on a slide than it does at the meter.

These two problems share a root cause. Both water use and carbon offsets are treated as reporting exercises rather than engineering targets. The water line is left out of the report, and the carbon line is papered over with credits. An infrastructure that genuinely reduces its footprint has to attack both at the source, in the design of the cooling system and the choice of generation, rather than in the accounting.

The indirect water figure is the part most reports never surface, and it is the larger number by far. The roughly 211 billion gallons embedded in the electricity that data centers consumed in 2023 dwarfs the 17.4 billion gallons they drew directly, a ratio of about twelve to one (Source 1). That water is consumed upstream at thermoelectric power plants that use evaporative cooling of their own, which means a data center that buys grid power inherits a water footprint it never measures and cannot control. Reducing the direct draw with efficient cooling is necessary but not sufficient, because the far larger indirect draw is a function of how the electricity itself is generated. A genuinely low water footprint therefore requires attention to both the cooling method at the facility and the generation mix behind the meter, which is exactly the pairing that offset driven accounting ignores.

Real-World Example

The public figures make the water problem concrete. With United States data centers directly consuming 17.4 billion gallons in 2023 and about twelve times that indirectly (Source 1), and with a single large facility using up to 5 million gallons a day during the very heatwaves when communities face shortages (Source 2), water is no longer a footnote. It is the resource line that determines whether a facility can be sited at all.

Consider an illustrative scenario. A regional operator plans a dense AI cluster in an arid, low cost power region, the kind of place where compute economics are best but water is scarcest. A conventional evaporative cooled design would put the facility in direct competition with the community for water during summer peaks. This scenario is illustrative and does not describe a specific Flux Core customer. Instead the operator deploys Flux Core units whose closed loop cooling uses zero outside water, removing the conflict entirely and making a site that would otherwise be off limits viable.

How Flux Core Solves It

Flux Core engineers the water line to zero at the source. The closed loop liquid cooling system recirculates a single fluid with no cooling tower and no evaporative loss, so there is no make up water permit and no draw on local supply. That directly answers the 17.4 billion gallon direct consumption problem and removes the facility from competition with communities during heatwaves (Source 1, Source 2). The approach is not a bet on unproven technology. NVIDIA has introduced a closed loop liquid cooling design that reduces on site water toward zero, independent validation that water free cooling is the direction the industry is moving (Source 3).

Carbon is engineered down rather than offset. Where a site allows, Nexus units run power agnostic on renewable solar paired with battery energy storage, displacing fossil power at the source rather than buying a credit against it. Where gas is needed, Flux Core prioritizes stranded and associated gas that would otherwise be flared or vented, so the project reduces emissions that are already happening instead of adding new load to a strained grid. The residual footprint is then addressed with avoided flaring value and carbon mechanisms matched against real, physical reductions.

Cooling finishes the carbon job as well as the water job. By removing the parasitic draw of air handling and the evaporative losses of cooling towers, closed loop liquid cooling holds efficiency high so more of every watt reaches the GPUs rather than the cooling plant. Delivered as a 1MW containerized Nexus R1000 built to Tier III design standards, the platform makes renewable first generation, flare remediation, and water free cooling built in rather than bolted on. That is a footprint claim that holds up when an auditor looks past the credits to the meter and the water line.

Crucially, the power agnostic architecture also attacks the indirect water problem that most operators never even measure. Because Flux Core can generate power on site from renewables or from gas that would otherwise be flared, the deployment does not draw its electricity from a grid whose thermoelectric plants consume roughly twelve times more water than the data center does directly (Source 1). Owning the generation path lets Flux Core shrink both the direct and the indirect water footprint at once, which no amount of renewable energy credit purchasing can accomplish for a grid connected, evaporative cooled facility. When the water line is engineered to zero at the facility and the generation is chosen to avoid water intensive upstream power, the resulting footprint is not a reporting artifact. It is a physical fact that survives the kind of scrutiny communities and regulators are increasingly bringing during heatwaves and droughts (Source 2).

The siting consequences are as important as the sustainability ones. A facility that needs no outside water is not forced onto the scarce, expensive parcels near abundant water and grid capacity. It can go where the cleanest or cheapest power is, including arid regions with strong solar resources or basins with stranded gas, precisely the places a water hungry conventional design cannot serve. In that sense the water free architecture is not only an environmental feature. It is what makes renewable first and flare reducing generation practical in the first place, because it removes the water constraint that would otherwise dictate where the compute has to sit.

Conclusion and Call to Action

When a single AI data center can drink 5 million gallons a day and the sector's indirect water runs twelve times its direct use (Source 1), and when demand peaks during the heatwaves communities can least afford (Source 2), water is the honest measure of a footprint. Flux Core engineers it to zero with closed loop cooling now validated across the industry (Source 3), paired with renewable first, flare reducing generation. If you want infrastructure that survives a real water and carbon audit, ask Flux Core to walk the full stack. Own the Energy, Monetize the Data.

Sources

1. ITIF, analysis of the data center water problem citing LBNL consumption figures. https://itif.org/publications/2026/07/06/the-data-center-water-problem-is-soluble/

2. Al Jazeera, United States heatwave and AI data center energy and water demand. https://www.aljazeera.com/economy/2026/7/3/us-heatwave-raises-alarms-over-ai-data-centre-energy-demands

3. Business Today, NVIDIA closed loop liquid cooling reducing water and energy use. https://www.businesstoday.in/technology/news/story/nvidia-says-its-new-liquid-cooling-system-can-reduce-water-and-energy-use-for-ai-data-centre-538735-2026-06-23

6. Sovereign AI, Sovereign Data: When Nations Fund Their Own Compute

Keywords: sovereign AI infrastructure, data residency compute, national AI capacity, air gapped GPU cluster, secure enclave data center

Executive Summary

Sovereign AI has moved from policy statement to funded budget line, and the numbers are now large enough to reshape where compute gets built. The global sovereign AI infrastructure market is estimated at about 24.8 billion dollars in 2026 and projected to reach 301.6 billion dollars by 2040, with spending expected to surpass 100 billion dollars in 2026 alone (Source 1). National commitments confirm the shift: the United Kingdom has pledged over 1 billion pounds to sovereign compute, and Israel has committed 30 billion dollars to national AI sovereignty (Source 2), while private capital such as Adani's announced 100 billion dollars in renewable powered AI data centers points the same direction. Sovereignty is no longer aspirational. It is capitalized. Flux Core Data Systems and Energy places dedicated GPU compute on the customer's own ground, inside a perimeter they control end to end, powered by a power agnostic energy path and cooled with zero outside water. This paper explains why funded sovereignty changes the buying decision and how Flux Core delivers compute that stays inside the perimeter.

The instrument of that control is the Nexus R1000, a 1MW containerized data center that arrives as a self contained unit with its own power path and its own closed loop cooling that uses zero outside water. Because it is power agnostic, it can operate in a secure facility, a research enclave, or a remote site without depending on a public grid or a municipal water supply, and because it is relocatable and deployable in three to six months, sovereign capacity can be stood up and scaled on the mission owner's timeline. Built to Tier III design standards, it lets a nation own the compute, the data, and the power as one controlled system.

The Problem

For years, sovereign AI was a principle without a budget. Governments spoke about keeping citizen data and national models inside their borders, but the compute to do so still largely lived in a handful of hyperscale clouds concentrated in a few countries. That is changing fast. The sovereign AI infrastructure market is now about 24.8 billion dollars in 2026 and is projected to grow to 301.6 billion dollars by 2040, with more than 100 billion dollars of spending expected in 2026 alone (Source 1). When a market grows at that rate, the constraint stops being political will and becomes physical delivery.

The national commitments are concrete and large. The United Kingdom has pledged over 1 billion pounds to sovereign compute, and Israel has committed 30 billion dollars to national AI sovereignty (Source 2). Private capital is moving in parallel, with announcements such as 100 billion dollars in renewable powered AI data centers signaling that the money now exists to build sovereign capacity at scale. The question every government and institution faces is no longer whether to fund sovereign compute, but how to deploy it quickly and securely.

Public hyperscale cannot answer that need. Placing sovereign workloads in a shared cloud, even a regional zone, means routing sensitive data through infrastructure the customer does not own and cannot fully control. Residency, classification, and air gap requirements cannot be met by a contractual assurance alone. A nation that has committed billions to sovereignty cannot have that sovereignty depend on another party’s terms of service.

Delivery is the binding constraint. Sovereign programs frequently need compute in secure facilities, research enclaves, or forward locations where grid power is limited or unreliable and where water for conventional cooling is scarce. Conventional construction on a multi year timeline, dependent on a fragile utility connection, cannot match the pace at which this funded demand is arriving. Sovereignty that cannot be deployed on time is sovereignty in name only.

The concentration problem gives sovereignty its urgency. The analysis behind the market forecast frames sovereign AI as a theme that cannot afford to concentrate in a single region or a handful of hyperscale providers, because doing so would leave nations dependent on infrastructure and jurisdictions they do not control (Source 1). A country that runs its national models and citizen data through a foreign cloud has outsourced not just its compute but its leverage. That is why the commitments are so large and so explicitly national: the United Kingdom’s billion pound pledge and Israel’s 30 billion dollar commitment are not merely spending, they are declarations that critical AI capacity must sit under domestic control (Source 2). The physical form of that control is dedicated hardware on sovereign ground, powered and cooled independently, which is a fundamentally different requirement from a contract with a public cloud region.

Real-World Example

The capital wave is the evidence. A sovereign market projected past 100 billion dollars in 2026 spending and toward 301.6 billion dollars by 2040 (Source 1), backed by the United Kingdom’s billion pound pledge and Israel’s 30 billion dollar commitment (Source 2), describes a level of funded demand that will require physical infrastructure delivered quickly and under the customer’s control.

Consider an illustrative scenario. A national research agency has funding to build sovereign AI capacity but needs it inside a secure facility where a public cloud region is not an acceptable answer and grid power is constrained. This scenario is illustrative and does not describe a specific Flux Core customer. The agency deploys self contained Flux Core units on its own ground, standing up dedicated GPU compute inside a perimeter it controls, without routing any workload through a shared cloud. The deployment meets residency and air gap requirements through its architecture rather than through a vendor’s promise.

How Flux Core Solves It

Flux Core places dedicated GPU compute on the customer’s own ground, inside a perimeter the customer owns end to end. Each Nexus unit is self contained and deployable, so sovereign compute can stand up in a secure facility, a research enclave, or a forward austere site without routing sensitive workloads through a shared cloud. Residency, classification, and air gap requirements are satisfied by the architecture itself, not by contractual assurances, which is exactly what a government committing billions to sovereignty needs (Source 1, Source 2).

Data independence rests on power independence. Because the platform is power agnostic, it runs on stranded or associated gas, on renewable solar paired with battery energy storage, on a hybrid microgrid, or on local grid power, so the mission never depends on a fragile utility tie. A sovereign deployment in a location with limited grid access is not blocked by an interconnection queue, because Flux Core brings its own power path. That is what lets funded sovereignty become deployed sovereignty on the timeline the capital demands.

Closed loop liquid cooling with zero outside water lets dense clusters run where water and infrastructure are scarce, including secure and remote sites that could never host a water hungry conventional facility. The Nexus R1000 delivers this as a 1MW containerized data center, built to Tier III design standards targeting 99.982 percent uptime, deployable in three to six months. Sovereignty then means what it should: owning the compute, the data, and the power as a single controlled system, rather than renting access to someone else's cloud and trusting their assurances.

The containerized form factor is what makes funded sovereignty deployable on the timeline the capital demands. A market spending more than 100 billion dollars in 2026 alone (Source 1) cannot wait years for conventional construction, and a sovereign program cannot compromise its perimeter to gain speed. Flux Core resolves that tension by delivering the perimeter, the power, and the cooling as a single self contained unit that arrives in months and can be sited where the mission requires rather than where the grid and the water happen to be. As requirements grow, units are added in parallel, so capacity scales without re architecting the security model. That combination, sovereign control plus rapid, repeatable delivery, is precisely what a nation that has committed billions to its own AI future needs in order to turn a budget line into operational compute.

The model scales the way sovereign programs actually grow. A nation can begin with a single unit for a pilot or a classified workload and add units as its needs and its budget expand, without ever moving data outside the perimeter or renegotiating a cloud contract. Each addition inherits the same security posture, the same independent power path, and the same water free cooling, so growth does not dilute control. That is the practical difference between sovereignty as a policy statement and sovereignty as operational infrastructure, and it is why a self contained, power agnostic platform is the natural home for the capital that nations are now committing to their own compute (Source 1, Source 2).

Conclusion and Call to Action

With the sovereign AI market passing 100 billion dollars in 2026 spending and heading toward 301.6 billion dollars by 2040 (Source 1), and with national commitments like the United Kingdom's billion pounds and Israel's 30 billion dollars now on the books (Source 2), sovereignty is funded policy that needs physical delivery. Flux Core delivers it: dedicated compute on the customer's ground, power agnostic and cooled with zero outside water, inside a perimeter the customer controls. If you are standing up national or mission owned AI, talk to Flux Core about sovereign compute on your terms. Own the Energy, Monetize the Data.

Sources

1. Seedscope, sovereign AI market size and the trillion dollar theme. https://seedscope.ai/blog/sovereign-ai-the-one-trillion-dollar-theme-that-cannot-concentrate-in-silicon-valley

2. Let’s Data Science, Israel commits 30 billion dollars to national AI sovereignty. https://letsdatascience.com/news/israel-commits-30b-to-national-ai-sovereignty-4215c922

7. College and University Data: Research That Cannot Wait in the Queue

Keywords: campus GPU cluster, university research computing, academic HPC capacity, on-campus AI training, research data sovereignty

Executive Summary

A research grant runs on a calendar, but a grid connection runs on a queue, and the two no longer fit together. In the PJM market, data centers now account for 94 percent of peak load growth through 2030, and during the July heatwave day ahead power prices leapt from 44 dollars to 436 dollars per megawatt hour, roughly a 900 percent move, with a one hour peak above 1,200 dollars (Source 1). A campus that needs to expand its power to host a serious AI cluster is entering that same crowded, volatile environment, where a substation upgrade can outlast the grant that justified it. Flux Core Data Systems and Energy brings dedicated research compute directly to campus without waiting on a multi year power upgrade, delivering up to 1MW of GPU capacity in a containerized unit that is power agnostic and cooled with zero outside water. This paper explains why the grid queue is now a research bottleneck and how Flux Core lets institutions run the work that cannot wait.

The campus deployment centers on the Nexus R1000, a 1MW containerized data center that delivers dedicated GPU capacity for high performance computing and model training on institutional ground. Its closed loop liquid cooling uses zero outside water, so it can sit on a campus without a new water permit, and its power agnostic design lets it draw on existing capacity, on solar with battery storage, or on on site generation instead of waiting for a substation upgrade. Built to Tier III design standards and live in three to six months, it gives research directors owned, dependable capacity on a grant timeline rather than a utility one.

The Problem

The grid that campuses depend on is being consumed by a single category of customer. In PJM, the largest grid market in the country, data centers now account for 94 percent of peak load growth through 2030, and the system is expected to grow from roughly 180,000 megawatts to more than 220,000 megawatts within a decade (Source 1). When one customer class drives almost all the growth, everyone else, including universities, waits behind it for capacity and for connection.

The volatility is as dangerous as the wait. During the July heatwave, PJM day ahead power ran from 44 dollars to 436 dollars per megawatt hour, close to a 900 percent increase, with a one hour peak above 1,200 dollars, and the reliability strain was severe enough that NERC issued a rare alert after data centers spontaneously severed grid connections (Source 1). A research institution that expands its grid dependence to power a large cluster inherits that price volatility and that reliability risk, neither of which fits a fixed research budget.

The cloud alternative fails research on its own terms. Shared cloud credits run out in the middle of a long training run, and a grant will not fund a rerun at on demand rates. Sensitive datasets, work governed by an institutional review board, and grant deliverables all sit more comfortably on infrastructure the institution controls than on a metered public cloud. Yet building on campus traditionally meant a multi year power upgrade that no grant timeline can absorb.

So research that cannot wait ends up waiting anyway. The science is ready, the funding is in place, and the faculty are hired, but the power to run the compute is stuck behind the same data center load that now dominates grid growth (Source 1). The institution is caught between a queue it cannot skip, a cloud bill it cannot sustain, and a grant clock that does not stop.

The scale of the grid transition makes the wait structural rather than temporary. PJM alone expects to grow from roughly 180,000 megawatts to more than 220,000 megawatts within a decade, and SemiAnalysis notes that verified grid additions run only about 15 gigawatts a year nationally while incremental data center demand climbs steeply, with more than half of new data centers expected to adopt behind the meter power after 2028 (Source 1, Source 2). For a university, that means the grid path is not merely slow today and fast tomorrow. It is congested for the foreseeable future, because the additions the grid can verify each year are already spoken for by the data center load ahead of it in line. An institution that ties its research computing roadmap to a utility upgrade is tying it to a queue that the data itself says will not clear on an academic timeline.

Real-World Example

The market data frames the squeeze. With data centers driving 94 percent of PJM peak load growth and heatwave prices spiking roughly 900 percent (Source 1), a campus trying to add substantial power for research computing is entering a market that is both crowded and volatile. The grid path is neither fast enough nor stable enough for a fixed grant budget.

Consider an illustrative scenario. A university has won a multi year grant that funds a large model training program, but its campus substation cannot supply the additional load and the utility upgrade would take longer than the grant period. This scenario is illustrative and does not describe a specific Flux Core customer. The university sites a containerized Flux Core deployment on institutional ground, delivering up to 1MW of GPU capacity live in months, and keeps sovereign control of its sensitive datasets and grant deliverables rather than metering hours against a public cloud bill.

How Flux Core Solves It

Flux Core brings dedicated research compute directly to campus without waiting on a multi year power upgrade. A containerized Nexus deployment delivers up to 1MW of GPU capacity for high performance computing, model training, and research computing, sited on institutional ground and live in three to six months. Research computing directors keep sovereign control of sensitive datasets, work governed by an institutional review board, and grant deliverables, rather than renting metered hours from a public cloud whose costs can consume a grant.

The deployment fits how campuses actually operate. Because the platform is power agnostic, the cluster can draw on existing grid capacity where it is available, on solar paired with battery energy storage, or on on site generation, which lets the institution avoid the crowded interconnection queue and the price volatility that the PJM data documents (Source 1). Aligning with tight budgets and campus sustainability pledges becomes practical rather than aspirational, because the energy path can be matched to the institution’s resources and commitments.

Closed loop liquid cooling with zero outside water keeps utility costs and environmental reporting clean, and it lets the deployment sit on a campus without the water permitting a conventional cooled facility would require. The Nexus R1000 is built to Tier III design standards targeting 99.982 percent uptime, so a training run is not interrupted by the reliability strain that pushed NERC to issue an alert on the public grid (Source 1). The institution gets dependable, dedicated capacity on a timeline that matches the grant, not the queue, which is exactly what research that cannot wait requires.

There is a governance dimension that matters as much as the power one. Research computing directors are accountable for sensitive datasets, for work governed by an institutional review board, and for the security commitments attached to federal and industry grants. A dedicated on campus deployment keeps all of that inside the institution’s own perimeter, under its own access controls, rather than distributed across a public cloud whose shared tenancy and metered pricing sit uneasily with both compliance requirements and fixed budgets. Because the platform is power agnostic and does not depend on the crowded grid that PJM data describes (Source 1), the institution can plan its research capacity with the same confidence it plans its laboratory space, as an owned asset with a known cost rather than a variable bill exposed to the price spikes that the July heatwave produced. That predictability is what lets a multi year research program actually run to completion on the schedule its funding assumes.

The result reframes what campus computing can be. Instead of choosing between a public cloud that meters every hour and a grid upgrade that outlasts the grant, an institution owns dedicated capacity that behaves like laboratory infrastructure, planned once and used for years. Faculty recruit against it, grants are written around it, and sensitive work stays inside the institution’s own perimeter. Because the deployment is power agnostic and cooled with zero outside water, it fits both the budget and the sustainability commitments that campuses increasingly must honor, turning research that would have waited in a queue into research that runs on schedule.

Conclusion and Call to Action

When data centers drive 94 percent of PJM peak load growth and heatwave prices spike roughly 900 percent (Source 1), a campus cannot afford to put its research on a multi year grid queue. Flux Core delivers up to 1MW of dedicated GPU capacity on campus in months, power agnostic and cooled with zero outside water, under the institution’s own control. If your research ambitions are outrunning your campus power and the grid queue, ask Flux Core about dedicated university AI infrastructure. Own the Energy, Monetize the Data.

Sources

1. Matter Fact, PJM data centers at 94 percent of peak load growth and heatwave price spikes. https://www.matterfact.com/newsletter/2026-07-07-data-centers-pjm-94-percent

2. SemiAnalysis, United States grid constraints and behind the meter adoption. https://newsletter.semianalysis.com/p/us-grid-constraints-towards-40gw

8. Powered Land Lease Rates: Renting a Bottleneck That Is Disappearing

Keywords: powered land lease rates, behind the meter power, data center ground lease, on-site generation economics, powered land bubble

Executive Summary

A powered land lease is a bet that grid scarcity will last, and that bet is getting worse by the quarter. A recent quote put 4MW of powered land at 3 million dollars a year, which works out to 750,000 dollars per megawatt annually just to hold land and a queue position, with electricity billed separately on top and before a single GPU is installed. Meanwhile the scarcity that justifies that rent is eroding fast. There are now 59 announced behind the meter sites totaling roughly 90 GW, more than a quarter of planned United States data center capacity, with about 82 GW announced in 2025 or later (Source 1), and the behind the meter power market is projected to grow from 18.4 billion dollars in 2025 to 198.6 billion dollars by 2034, a 30.5 percent compound annual growth rate (Source 2). Flux Core Data Systems and Energy generates power rather than renting access to it, on a power agnostic platform cooled with zero outside water. This paper explains why the powered land premium is a melting asset and why owning the energy beats renting the bottleneck.

The alternative to the lease is the Nexus R1000, a 1MW containerized data center that Flux Core sites where the power resource is strongest rather than where a queue position happens to sit. With a power agnostic design spanning stranded gas, solar with battery storage, hybrid microgrids, and grid power, and with closed loop cooling that uses zero outside water, the unit is free of the water permits and interconnection dependence that tether conventional facilities to expensive parcels. Built to Tier III design standards and deployable in three to six months, it turns generation, not ground rent, into the foundation of the project.

The Problem

Powered land commands a premium only because grid connections are scarce and slow. The recent quote of 750,000 dollars per megawatt per year for 4MW, with electricity billed separately, is a price for scarcity, not for value created. The tenant pays 3 million dollars a year to occupy land and hold an interconnection position before any compute is installed or any token is served. That rent is sustainable only as long as the grid remains the sole gateway to power.

That gateway is opening. There are now 59 announced behind the meter sites totaling about 90 GW, representing more than a quarter of planned United States data center capacity, with roughly 82 GW of that announced in 2025 or later (Source 1). Behind the meter generation is no longer a niche workaround. It is becoming a mainstream way to power data centers, and each site that generates its own power is a site that does not need a powered land lease.

The market forecasts confirm the direction. The behind the meter power generation market is projected to grow from 18.4 billion dollars in 2025 to 198.6 billion dollars by 2034, a compound annual growth rate of 30.5 percent (Source 2). SemiAnalysis expects that after 2028 more than half of new United States data centers will adopt behind the meter power (Source 3). Capital is backing the shift at scale, with Brookfield and Bloom Energy expanding on site fuel cell financing to 25 billion dollars for up to 1 GW of capacity (Source 4). When self generated power becomes the default, the exclusive access to grid power that a powered land lease sells loses its value.

This is the anatomy of a commercial real estate bubble. Landlords are writing decade long leases priced against a scarcity that a competing technology is on track to eliminate within the first years of those leases. The demand for AI compute is real and growing, but the assumption that the only way to serve it is through grid connected land is not. When that assumption breaks, the leases do not gracefully reprice. They strand, because the value proposition behind them has evaporated.

It is worth being precise about why this is a real estate bubble and not an AI bubble, because the two are constantly conflated. Demand for AI compute is real, large, and growing, as the same data center load driving grid congestion demonstrates. That demand is not speculative. What is speculative is the assumption that the only way to serve it is through grid connected land, and therefore that grid connected land commands a permanent premium. Bubbles form when an asset is priced on an assumption that a structural change is about to invalidate. Powered land is priced on the assumption that the grid remains the sole gateway to power. Behind the meter generation, now 90 GW announced and forecast to serve more than half of new data centers after 2028 (Source 1, Source 3), removes that assumption, and when the assumption goes, the premium goes with it, even as the underlying AI demand keeps climbing.

Real-World Example

The quote itself is the example. At 750,000 dollars per megawatt per year for 4MW, a tenant pays 3 million dollars annually for land and a queue position, with electricity on top. Set that against the behind the meter trajectory: 90 GW already announced (Source 1), a market growing to 198.6 billion dollars by 2034 (Source 2), and more than half of new data centers expected to go behind the meter after 2028 (Source 3). Against that backdrop, a decade long powered land lease is a bet against the clear direction of the market.

Consider an illustrative scenario. A developer faced with that powered land quote runs the numbers and declines. Instead of paying 750,000 dollars per megawatt to rent a queue position, the developer generates power on site behind the meter and energizes in months. This scenario is illustrative and does not describe a specific Flux Core customer, but it is precisely the choice the market data says developers are increasingly making (Source 1, Source 3). Every developer who makes that choice removes a future tenant from the powered land market, which is exactly how a rent roll collapses.

How Flux Core Solves It

Flux Core stands on the other side of the powered land trade. We understand power, so we develop our own opportunities rather than rent someone else’s bottleneck. There is no ground lease to pay, no interconnection queue to wait in, and no electricity pass through markup layered on top of the rent. The 750,000 dollar per megawatt structure simply does not exist in the Flux Core model, because Flux Core is not renting scarcity. It is generating power.

The deployments are power agnostic across stranded or associated gas, renewable solar paired with battery energy storage, hybrid microgrids, and grid power where it is available. That flexibility is what makes the behind the meter strategy practical, and it aligns Flux Core with the shift that the announced 90 GW pipeline and the market forecast describe rather than against it (Source 1, Source 2, Source 3). Instead of betting on a scarcity that is eroding, Flux Core is helping erode it, in the same direction that the Brookfield and Bloom financing expansion points (Source 4).

The Nexus R1000 delivers this as a 1MW containerized data center, deployable in three to six months, built to Tier III design standards, and cooled by a closed loop liquid system that uses zero outside water. That last point matters for siting, because a deployment free of water permits and cooling towers is not tethered to the scarce, expensive, queue burdened parcels that command powered land rents. Flux Core can site where the power is, not where the queue position happens to be, which is the structural advantage that makes the powered land premium irrelevant to its business and to its customers.

Consider what the same 4MW looks like under the two models. Under the powered land lease, the tenant pays 3 million dollars a year for land and a queue spot, waits years for interconnection, and then pays for electricity on top with a pass through markup. Under the Flux Core model, four Nexus units are deployed on a site chosen for its power resource, energized in months, and cooled without a drop of outside water, with the energy owned and operated rather than rented at a marked up rate. The contrast is not incremental. One path pays a premium to wait in a queue that data center load has already crowded, while the other skips the queue entirely by generating power where the compute sits. As the behind the meter market grows toward 198.6 billion dollars by 2034 (Source 2) and self generation becomes the default, the first path strands and the second becomes the norm, which is why Flux Core has built its entire model on owning the energy rather than renting access to it.

The broader point is a matter of which side of the transition an organization wants to stand on. Renting powered land is a bet that the grid stays the only gateway to power, a bet that 90 GW of announced behind the meter capacity and a market on track for 198.6 billion dollars by 2034 are steadily undermining (Source 1, Source 2). Generating power is a bet that self generation becomes the norm, which is the direction every credible forecast points (Source 3). Flux Core has built its entire model on the second bet, so its customers are aligned with the market shift rather than exposed to it, and they capture the value that a landlord charging ground rent is about to lose.

Conclusion and Call to Action

A 750,000 dollar per megawatt powered land lease rents a bottleneck that 90 GW of announced behind the meter capacity (Source 1), a market growing to 198.6 billion dollars by 2034 (Source 2), and a majority behind the meter shift after 2028 (Source 3) are steadily dismantling. Flux Core owns the energy instead of renting the bottleneck, with a power agnostic, zero water platform and no ground lease, no queue, and no electricity markup. Stop renting a power problem that is about to reprice. Talk to Flux Core about owning the energy and monetizing the data. Own the Energy, Monetize the Data.

Sources

1. Matter Fact, 59 announced behind the meter sites totaling roughly 90 GW. https://www.matterfact.com/newsletter/2026-07-03-data-centers-done-waiting-grid

2. Research Intelo, behind the meter power generation market size and forecast. https://researchintelo.com/report/ai-data-center-behind-the-meter-power-generation-market

3. SemiAnalysis, more than half of new data centers behind the meter after 2028. https://newsletter.semianalysis.com/p/us-grid-constraints-towards-40gw

4. Construction AI Brief, Brookfield and Bloom Energy expand on site financing to 25 billion dollars. https://constructionaibrief.com/posts/2026-07-02-brookfield-bloom-national-grid-onsite-power-turbine-backlog

---