Flux Core Data Systems and Energy

1. Offtake Agreements Reggie

Compute without a buyer is just an expensive science project. What separates a real deployment from a slide deck is contracted demand, and that is where most operators fall short. Flux Core builds from committed offtake outward. We already hold a Memorandum of Understanding covering 50 to 100 MW of demand, and we are adding offtake partners to the book now.

The demand base is intentionally diversified so no single buyer sets the pace: neoclouds reselling GPU hours, enterprises hosting private models, universities and national labs, hospitals, municipalities, and AI platforms scaling inference. Every partner reserves capacity inside a containerized Nexus deployment that energizes in months, cooled by a closed loop liquid system that consumes zero outside water.

Power agnostic siting is what makes those commitments deliverable. Capacity can sit on stranded gas, on solar paired with battery storage, or on grid power, so an offtake is never trapped behind a multi year interconnection request.

Sourcing committed AI capacity, or ready to anchor a project as an offtake partner? Contact Flux Core to review the megawatts still open.

2. Bare Metal Provisioning (Powered Shell up to 1.3MW IT Load) Dylan

Cooling is where most powered shells quietly fail an AI tenant. A shell can offer power and four walls and still choke the moment you rack 100 kW cabinets. Flux Core solves the thermal problem before you ever ship hardware, delivering a bare metal environment provisioned for up to 1.3MW of IT load and validated for the densest GPU clusters.

The loop is closed and water free. Our closed loop liquid cooling recirculates the same fluid and draws zero outside water, so there is no cooling tower, no make up water permit, and no evaporative loss capping your density. Racks run hot chip dense at Tier III reliability, and you bring your servers or lease ours into a space that is ready on day one.

The shell is power agnostic. Site it on stranded natural gas, on renewable solar with battery storage, or on an existing grid tie, and Flux Core builds the power path to match, which is how a 1.3MW environment reaches edge and power constrained locations traditional colocation cannot serve.

Racking high density GPUs? Ask Flux Core about a thermally proven powered shell up to 1.3MW.

3. Power as a Service KP

Power is only ten to fifteen percent of a data center's total cost, yet it is the single factor that decides whether the facility ever earns a dollar. Flux Core turns that binding constraint into a simple purchase. With Power as a Service, we finance, own, and operate the generation and distribution, and you buy reliable power for your compute as an operating expense.

No capital tied up in generators, switchgear, or batteries. No years lost to a utility interconnection study. Because Flux Core is power agnostic, the system beneath your load can be stranded natural gas, solar paired with battery storage, a hybrid microgrid, or a grid tie, sized to your ramp and your site.

Efficiency stays intact on the thermal side. Closed loop liquid cooling with zero outside water holds PUE low and keeps your facility off local water supply, which matters most in the arid, low cost power regions where compute increasingly wants to live. You pay a predictable rate while Flux Core captures the credits and tax structures underneath.

Tired of your project living or dying by the grid? Ask Flux Core how Power as a Service gets you online.

4. Oil and Gas Well Deployments Reggie

For decades a well with no pipeline meant gas sold at a discount or burned at the flare. Flux Core rewrites that outcome. We bring containerized data centers to the wellhead and convert stranded and associated gas directly into AI compute revenue, on the pad, with no midstream takeaway required.

The value proposition for an operator is immediate. Gas that is currently flared, vented, or moved at a steep basis discount becomes fuel for on site generation that powers high margin compute and token output. It is a fresh revenue line off an existing asset, and cutting the flare improves the emissions profile in the same move.

In a basin, water is the scarce resource, so our closed loop liquid cooling answers it by design, using zero outside water. The unit trucks to a remote location and runs power agnostic on wellhead gas today, ready to shift to renewables or grid power as conditions change, and it relocates as the field's production curve moves.

Sitting on gas with nowhere to go? Flux Core turns it into compute. Let's scope a wellhead deployment on your acreage.

5. Zero Carbon Footprint Dylan

Most operators buy their way to carbon neutral with unbundled offsets after the fact. Flux Core builds the reductions into the deployment itself, and here is the actual mechanism behind each step.

Start with generation. Where the site supports it, Nexus units run on renewable solar paired with battery energy storage, displacing fossil power at the source. Where gas is required, we prioritize stranded and associated gas that would otherwise be flared or vented, so the project cuts emissions that were already occurring rather than creating new ones.

Next, close the remaining gap. Avoided flaring value and carbon credit mechanisms are matched against the residual footprint to drive net emissions toward zero. Then cooling does the rest: a closed loop liquid system using zero outside water removes evaporative loss and the parasitic energy of air cooling, holding PUE low so more of every watt reaches the GPUs.

Renewable first generation, flare remediation, carbon credits, and water free cooling, engineered together rather than offset afterward. That is a defensible path to zero carbon.

Building infrastructure that has to pass a real ESG review? Ask Flux Core to walk the carbon stack.

6. Sovereign AI, Sovereign Data KP

The fastest growing requirement in AI is not more FLOPs. It is control: the guarantee that a model and its data never leave a defined perimeter or jurisdiction. Public hyperscale cannot make that promise, and that gap is exactly where Flux Core operates. We site dedicated GPU compute on the customer's own ground, inside a boundary they own 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 location without shipping sensitive workloads to a shared cloud. Data residency, classification, and air gap requirements are met by architecture, not by a vendor's policy language.

Power independence underwrites data independence. Because the platform is power agnostic, it runs on stranded gas, on solar with battery storage, or on local grid power, so the mission is never hostage to a fragile utility tie. Closed loop liquid cooling with zero outside water lets dense clusters operate where water and infrastructure are scarce.

Real sovereignty means owning the compute, the data, and the power together. Talk to Flux Core about deploying sovereign AI on your terms.

7. College and University Data Reggie

Faculty are being recruited on the promise of AI research the campus infrastructure cannot actually support. The GPUs are budgeted, the grant is awarded, and then the substation and a multi year utility upgrade stand in the way. Flux Core clears that blocker, putting dedicated research compute on campus without waiting on a capital power project.

A containerized Nexus deployment delivers up to 1MW of GPU capacity for HPC, model training, and research computing, sited on institutional ground and live in months. Research computing directors retain sovereign control of sensitive datasets, IRB governed work, and grant deliverables instead of renting hours in a shared public cloud.

It also matches how institutions run. Power agnostic means the cluster can draw on existing grid capacity, on solar with battery storage, or on site generation, aligning with tight budgets and campus sustainability commitments alike. Closed loop liquid cooling with zero outside water keeps utility costs and environmental reporting clean.

Have research ambitions outrunning campus power? Ask Flux Core about dedicated university AI infrastructure.

8. Powered Land Lease Rates — The Real Bubble Dylan

A quote just crossed our desk: 4MW of powered land for 3 million dollars a year, with electricity billed on top. That is 750,000 dollars per megawatt, per year, for dirt and an interconnection promise. The compute is not in the building. The chips are not in the rack. You are renting the right to wait in a five year grid queue and paying a premium for the privilege.

This is the part the market keeps getting wrong. The AI buildout is not the bubble. The bubble is powered land as a service. Landlords are underwriting decade long leases against scarcity that is already eroding, and when behind the meter generation becomes the default, those rent rolls collapse. Behind the meter capacity in the US is projected to exceed 40GW by 2028. The moment a developer can generate power on site, a 750,000 dollar per MW ground lease is worth a fraction of the paper it is printed on. That is not an AI correction. That is a commercial real estate bubble, and it will pop.

Flux Core sits on the other side of that trade. We understand power, so we develop our own opportunities instead of renting someone else's bottleneck. Our deployments are power agnostic, running on stranded gas, renewable solar paired with battery storage, or grid power, and cooled by a closed loop liquid system that uses zero outside water. No ground lease. No queue. No pass through markup on electricity.

Stop renting a power problem. Talk to Flux Core about owning the energy and monetizing the data.

White Paper: The New Architecture of AI Compute Offtake

Flux Core Data Systems1. Executive Summary: Beyond the Slide Deck

The AI infrastructure boom has revealed a fundamental mismatch between the speed of capital deployment and the reality of energy procurement. Traditional data center models—relying on multi-year interconnection queues and massive, centralized facilities—are failing to meet the agility requirements of modern AI developers. Compute without a buyer is an expensive science project; compute without a secure, flexible energy source is a stranded asset. Flux Core’s offtake model bridges this gap by decoupling high-density AI compute from the grid-dependent, long-lead-time constraints that stifle traditional operators.2. The Problem: The Interconnection Bottleneck

Most operators approach compute deployment as a real estate play, seeking to secure a grid connection first and building a shell around it. This “grid-first” strategy forces the developer into a multi-year queue, leaving the tenant hostage to utility timelines and escalating costs. The result is often “powered land” that is expensive, distant from the workload, and inflexible. When a developer cannot guarantee power availability, they cannot guarantee compute uptime.3. The Solution: Offtake-Led Deployment

Flux Core builds from committed demand outward. Our approach starts with a clear, financeable offtake agreement. By engaging with neoclouds, enterprises, and research institutions early, we validate the demand before a single container is deployed.

Real-World Application: The Red Oak Deployment

In our Red Oak, Iowa project, we moved from concept to site-ready status by ignoring the traditional grid-queue dependency. Instead of waiting for a multi-year utility upgrade, we engineered a containerized solution (10 racks, 256 NVIDIA B200 GPUs) that leverages 2x 720 kW natural gas generators in an N+1 paralleled configuration as the primary power source. The utility feed acts as a standby backup, not the critical path. This offtake-led strategy allowed us to move to a 6-month go-live timeline, ensuring our partners are not renting the right to wait in a queue, but buying the guarantee of active compute.4. Diversified Demand: The Foundation of Stability

A financeable data center requires a diversified demand base to mitigate risk. Flux Core's offtake strategy prioritizes a mix of:

- **Neoclouds/GPU Resellers:** Scaling inference capabilities on demand.
- **Enterprises:** Hosting private, high-security models.
- **Academic/Research:** Grant-funded HPC and model training.

By spreading demand across these segments, no single partner sets the pace, and the infrastructure is insulated from the volatility of any one sector. Each partner reserves capacity inside our containerized Nexus units, ensuring dedicated compute density.5. Conclusion: Power Agnostic, Mission Critical

The future of AI infrastructure is not about finding the biggest grid tie. It is about deploying compute where the power can be generated on-site, behind the meter. By matching specific GPU compute requirements (such as our B200-ready containers) with agile, financeable offtake agreements, Flux Core provides a path to immediate deployment.

We are currently expanding our offtake book for projects like Red Oak. Contact the Flux Core team to review the megawatts currently open for partnership.

White Paper: The Bare Metal Imperative – Engineering High-Density AI Environments

Flux Core Data Systems1. Executive Summary: The End of Legacy Colocation

The rapid acceleration of generative AI has created a fundamental infrastructure crisis. Traditional colocation facilities, built for the CPU-dense, low-wattage era, are reaching a thermal wall. When a

data center provider attempts to retrofit legacy air-cooled halls to accommodate the massive heat output of Blackwell-class GPU clusters, the result is either restricted performance or catastrophic failure.

Flux Core has pioneered a new model: "Bare Metal Provisioning." We do not just lease space; we deliver a thermally validated, 1.3MW-capacity environment engineered specifically for high-density AI clusters. By decoupling the hardware lifecycle from the real estate development cycle, Flux Core provides a path to immediate deployment, supporting NVIDIA B200, H200, and GB300 platforms at the edge.2. The Thermal Bottleneck[1]

The primary constraint in AI deployment is no longer just power; it is the ability to reject heat at scale. Air-cooled data centers generally fail to support cabinets exceeding 20–30 kW. As GPU clusters move toward 130 kW per rack, air-based cooling becomes spatially inefficient and energetically unsustainable.[2]

Most "powered shells" currently on the market are merely cold, empty boxes. They lack the dedicated thermal infrastructure required to keep high-density chips from thermal throttling. Flux Core addresses this by building the cooling solution *into* the envelope. Our proprietary closed-loop liquid cooling recirculates coolant through rear-door heat exchangers or direct-to-chip manifolds, rejecting heat to the atmosphere via dry coolers rather than cooling towers. This architecture serves two purposes:[2]

- **Performance:** It maintains consistent silicon temperatures, allowing for sustained training loads without throttling.
- **Resource Efficiency:** It eliminates the reliance on outside water, effectively removing a major permitting and supply-chain bottleneck that restricts hyperscalers and enterprises from deploying in arid or water-constrained regions.[1][2]

3. Bare Metal Provisioning: The "Day One" Architecture

Flux Core's Bare Metal offering is defined by readiness. We deliver a fully provisioned 1MW+ shell that includes power distribution, liquid cooling, and structured cabling. Customers simply rack their hardware—or lease ours—and begin compute operations immediately.[1]

This modular "Powered Shell" approach effectively commoditizes the infrastructure layer, allowing developers to focus on the compute stack rather than the building. Our technical specifications are optimized for next-gen accelerators:

- **Capacity:** Up to 1.3MW of IT load per shell.[2]
- **Distribution:** 480V three-phase distribution with 3000A busway and redundant PDUs, supporting high-density power delivery to the rack.[1]
- **Compliance:** Built to SOC 2 and NIST 800-171 standards, ensuring that dedicated bare metal is suitable for enterprise and sovereign workloads.[1]

4. Deployment Versatility: Power-Agnostic Infrastructure

A persistent issue with traditional bare metal hosting is the reliance on the local utility interconnection queue, which can force projects into years of delay. Flux Core's bare metal architecture is intentionally power-agnostic. Because our units are fully containerized and self-contained, they are not tethered to a static, grid-dependent substation.[1][2]

We leverage this flexibility to site infrastructure near the power source. As seen in our field operations (e.g., our West Texas deployments), we can deploy GPU clusters directly to wellheads, utilizing stranded natural gas that would otherwise be flared. This ability to site on-pad allows us to bypass the 5-year grid interconnection queue entirely, turning an environmental liability (flared gas) into an operational asset (compute revenue). Whether the input is renewable solar, stranded gas, or existing grid tie-ins, Flux Core ensures the shell is energized on the developer's timeline, not the utility's.5. Conclusion: A New Standard for AI Infrastructure[1][3][2]

The "powered shell" market is currently suffering from a bubble of optimism, where developers are charging massive premiums for empty dirt and vague interconnection promises. Flux Core rejects this model. We believe infrastructure should be judged by its ability to support the compute load, not by its proximity to a utility line.[2]

By delivering 1.3MW, thermally proven, liquid-cooled, and power-agnostic bare metal environments, Flux Core provides a tangible solution to the AI capacity gap. We are currently scoping projects for high-density inference and training workloads. Contact the Flux Core strategy team to evaluate your deployment requirements.

White Paper: Power as a Service – Unlocking Stranded Compute

Flux Core Data Systems1. Executive Summary: The OpEx Shift

In the conventional data center model, energy is treated as a construction problem—a multi-year capital investment characterized by switchgear, substations, and utility negotiations. For AI infrastructure, this approach is fundamentally broken. Power typically accounts for only 10% to 15% of a data center's total cost, yet it acts as the primary gatekeeper for the entire operation. If the power cannot be brought online, the compute revenue remains zero.

Flux Core's "Power as a Service" (PaaS) model reframes this dynamic. We decouple the energy system from the balance sheet, allowing developers and operators to treat power as a predictable operating expense (OpEx) rather than a multi-year capital drain. By financing, owning, and operating the generation assets, we eliminate the utility interconnection queue as a project risk, enabling compute to go live on a developer's timeline, not the utility's.2. The Capital Trap

The industry is currently plagued by the "interconnection bottleneck." Traditional operators are forced to commit millions of dollars in upfront capital to projects that may sit dormant for three to five years while waiting for grid upgrades. This capital-intensive model creates two critical failures:

- **Balance Sheet Exposure:** High upfront spending on generators, batteries, and switchgear reduces the capital available for GPU procurement and operational scaling.
- **Asset Stranding:** Projects are built where the grid exists, not where the power is cheapest or most available, leading to inflated operational costs and susceptibility to grid volatility.

Flux Core disrupts this by shifting the responsibility of the power infrastructure—from generation to distribution—to the provider. The developer is no longer a utility negotiator; they are a compute operator.3. The Solution: Dispatchable Power as an OpEx

Flux Core's PaaS platform is fundamentally **power-agnostic**. Our infrastructure is designed to integrate the most efficient energy source available at the specific site, whether that is stranded natural gas, solar paired with battery energy storage, or a hybrid microgrid.

By sizing the power system to the specific ramp-up needs of the client's GPU load, we ensure that energy expenditure scales linearly with compute revenue. This removes the "all-or-nothing" requirement of grid-tied deployments. Clients pay a predictable, contracted rate for reliable power, while Flux Core manages the volatility of fuel, storage, and grid-integration credits. This structure essentially turns the energy system from a complex technical burden into a utility-grade service.4. Real-World Application: Beyond the Grid

Traditional hyperscalers are often forced to choose between the grid and nothing. Our field experience demonstrates a third path. In our recent deployments, we have utilized behind-the-meter generation to power high-density clusters where grid infrastructure was non-existent.

By utilizing our proprietary closed-loop liquid cooling in tandem with PaaS, we create a high-efficiency stack that minimizes total energy consumption. Because our cooling system uses zero outside water, we can deploy PaaS solutions in arid, low-cost power regions—areas that were previously considered "unbuildable" by conventional data center operators due to water scarcity or grid limitations.

This integration of Power as a Service and advanced thermal management means that a project can be energized, provisioned, and optimized in a fraction of the time required by traditional methods.5. Conclusion: Energy Independence

The future of AI infrastructure belongs to those who control their power supply rather than those who rent it from a stagnant utility. Power as a Service is not just a financing model; it is a strategic advantage. It allows Flux Core partners to deploy in months, not years, and to maintain profitability regardless of local utility constraints.

We are currently evaluating new sites for Power as a Service integration. If your project is currently stalled by an interconnection study or you are seeking a move to an OpEx-based infrastructure model, contact the Flux Core strategy team to review your site's capacity potential.

White Paper: The Wellhead Advantage – Monetizing Stranded Assets for AI Compute

Flux Core Data Systems1. Executive Summary: Turning Liability into Logic

For decades, the oil and gas industry has faced a fundamental midstream constraint: when a well produces associated gas in a region without pipeline takeaway capacity, that gas becomes a liability. Historically, operators have been forced to flare this gas, vent it, or sell it at steep basis discounts—often negative—at hubs like Waha. This is not just a loss of revenue; it is an environmental and regulatory friction point.

Flux Core has re-engineered this paradigm. By deploying containerized, high-density AI data centers directly on the wellhead, we convert stranded and associated gas into a high-margin fuel source for GPU compute. This model transforms an operational waste stream into a direct revenue line, bypassing the midstream infrastructure trap entirely.2. The Problem: The Midstream Bottleneck

The current economics of oil production in basins like the Permian often create a "gas surplus" problem. As producers chase oil, they inadvertently produce associated gas that the local infrastructure cannot absorb.

- **The Flare Trap:** Regulations regarding flaring are tightening globally, increasing the risk of fines and ESG penalties.
- **Pipeline Scarcity:** Reliance on pipeline takeaway forces operators to accept whatever basis price the midstream market dictates—which is frequently sub-$3/MMBtu or worse.
- **Capping Liability:** When takeaway is full or costs are too high, operators are forced to throttle oil production, capping the profitability of the entire well pad.

3. The Solution: On-Pad Compute

Flux Core's Nexus deployment model brings the market to the molecule. By placing a containerized data center directly on the pad, we create an immediate "offtake" for the gas right at the wellhead.

The Operational Model:

Instead of moving gas to a pipeline, we move the compute to the gas. Our modular, skid-mounted units include all necessary compression and scrubbing technology required to clean the gas and feed it into our high-efficiency natural gas generators.

- **No Midstream Requirement:** By utilizing the gas on-site, the project is decoupled from pipeline capacity constraints.
- **Modular Scaling:** If a well's production curve shifts, the Nexus unit can be trucked to a new pad or relocated to match changing field output.
- **Integrated Infrastructure:** As discussed in our site engineering reviews, Flux Core handles the full integration—from the scrubber systems to the microgrid power distribution—ensuring that the operator's only responsibility is providing the gas stream.

4. Technical Integration and Reliability

Data centers require consistent power and thermal management. Our wellhead deployments utilize the same high-availability architecture found in our grid-tied facilities, tailored for the field environment:

- **Generation & Conditioning:** We deploy proprietary compression and scrubber systems to condition the wellhead gas, ensuring it meets the specification required for stable, industrial-grade power generation.
- **Thermal Management:** The Permian Basin and similar producing regions are notoriously water-constrained. Our closed-loop liquid cooling system is engineered for these environments. It uses zero outside water and relies on dry coolers, eliminating the need for cooling towers, make-up water permits, or the maintenance headaches of water-based thermal systems.
- **Real-Time Monitoring:** Through our all-encompassing Building Management System (BMS), including partnerships with Vast Storage and Lenovo, we provide comprehensive, real-time visibility into the health of the stack. This gives both Flux Core and the site operator the assurance that the compute remains online and the gas is being utilized at peak efficiency.

5. Conclusion: A No-Brainer for Production

For the producer, this is a clear strategic upgrade. Gas that is currently burned as waste is used to power high-value AI compute and token generation. It stabilizes the asset's economics, improves the emissions profile by eliminating the flare, and allows production to continue without reliance on midstream volatility.

We are currently expanding our portfolio of Permian and Mid-Continent wellhead projects. If you are an operator with stranded gas and acreage that could host a 1MW–3MW compute load, contact the Flux Core team to scope a wellhead deployment.

White Paper: The Engineered Path to Net-Zero – Sustainability as Hardware

Flux Core Data Systems1. Executive Summary: Beyond the Offsetting Game

Sustainability in the data center industry has historically been defined by financial engineering: purchasing unbundled Renewable Energy Certificates (RECs) or carbon offsets *after* a facility is built to claim "carbon neutral" status. This approach does nothing to alter the operational reality of the underlying infrastructure.

Flux Core rejects this performative model. We believe that true **net zero data center** operations must be engineered into the hardware and the energy architecture itself, not layered on top as a compliance expense. By focusing on **carbon-negative compute** potential, we build infrastructure that achieves radical efficiency through low **efficient PUE** (Power Usage Effectiveness), **water-free**

cooling, and the remediation of methane through **avoided flaring credits**.2. The Fallacy of Unbundled Offsets

Most modern operators are constrained by a fundamental design flaw: their dependence on the grid and air-cooled thermal systems. To offset the carbon intensity of an air-cooled, grid-tied facility, operators are forced to buy offsets indefinitely. This is a cost-heavy strategy that creates no long-term value and leaves the operator vulnerable to changing energy regulations and rising carbon pricing.

For enterprises and hyperscalers, this is no longer sufficient. Institutional investors, sovereign wealth funds, and ESG-conscious stakeholders require **ESG-ready infrastructure** that reduces the footprint at the source, rather than simply paying for someone else to reduce theirs.3. The Flux Core Architecture: A Three-Pillar Approach

Our path to a **net zero data center** is built on three technical pillars that maximize efficiency while minimizing the environmental tax of compute:

- **Pillar I: Renewable-First Generation**

  Where site conditions permit, we deploy **renewable microgrid** solutions—pairing solar or wind generation with utility-scale battery energy storage. This allows us to displace fossil-fuel grid consumption entirely.
- **Pillar II: Methane Remediation (Avoided Flaring)**

  Where on-site generation via natural gas is the optimal solution, we prioritize stranded and associated gas that would otherwise be vented or flared. By converting this gas into high-value compute output, we generate **avoided flaring credits**. We are not creating new emissions; we are capturing emissions that were already occurring and putting them to work.
- **Pillar III: Thermal Efficiency**

  The hidden cost of compute is water. Conventional cooling towers consume millions of gallons of water annually. Our proprietary closed-loop liquid cooling system uses **water-free cooling**, eliminating evaporative loss and drastically reducing the parasitic energy load of traditional air-conditioning units.

4. Real-World Application: Lowering the PUE

The physical reality of our commitment is visible in the metrics. While a standard legacy data center operates at a PUE between 1.4 and 1.6, our Nexus deployments are engineered to achieve a PUE of 1.2 to 1.3.

This 0.2 to 0.4 delta is not merely a rounding error; it is a fundamental shift in energy consumption. For every megawatt of compute, that reduction in PUE means hundreds of thousands of kilowatt-hours saved annually, directly reducing the total carbon load of the deployment. By combining this

efficiency with our flare-gas monetization strategy, we create a platform capable of achieving **carbon-negative compute** for our partners.5. Conclusion: Defensible ESG Infrastructure

The era of "buying" sustainability is ending. As regulatory bodies and global markets tighten their requirements for carbon disclosure, Flux Core provides the infrastructure that survives the audit. By stacking renewable-first generation, methane remediation, and water-free thermal design, we offer a path to net-zero that is not only defensible but also profitable.

We are currently scoping new sites for ESG-optimized deployment. If your firm is seeking to align its infrastructure with aggressive sustainability targets, contact the Flux Core strategy team to walk through our carbon stack.

White Paper: Sovereign AI and Sovereign Data – Infrastructure for a Controlled Perimeter

Flux Core Data Systems1. Executive Summary: The Sovereignty Requirement

The most critical requirement in the next phase of AI development is not just raw compute throughput—it is control. For government agencies, defense contractors, financial institutions, and critical infrastructure providers, the standard public cloud model presents an unacceptable risk. The guarantee that a model, its training data, and the resulting insights will never leave a defined jurisdiction or perimeter is now a non-negotiable prerequisite.

Flux Core redefines the architecture of control. We move away from the shared, multi-tenant cloud and deliver dedicated, self-contained AI infrastructure that sits entirely on the customer's ground. By owning the compute, the physical boundary, and the power generation, Flux Core allows organizations to maintain total data residency and operational air-gapping, meeting the most rigorous classification and security mandates.2. The Failure of the Shared Perimeter

Current hyperscale cloud providers operate on a shared infrastructure model. While they offer security protocols, the underlying reality is a multi-tenant environment where data is processed in facilities that are geographically distributed and governed by broad, vendor-controlled policies.

For sovereign-grade workloads, this architecture is insufficient. Organizations requiring true data sovereignty face three critical risks:

- **Jurisdictional Leakage:** Inability to guarantee exactly where data is processed or stored, creating compliance friction with national data sovereignty laws (e.g., GDPR, NIST, CMMC).
- **Logical Co-mingling:** The risk—however small—of logical co-mingling with non-secure or external datasets within a shared software-defined environment.
- **Perimeter Vulnerability:** Reliance on external connectivity to a public cloud, which creates an unnecessary attack surface and a single point of failure for the entire compute stack.3. The Solution: Hardened, Sovereign Nexus Deployments

Flux Core’s sovereign strategy is built on the Nexus unit: a self-contained, containerized GPU cluster that functions as a private, localized data center.

Architectural Sovereignty:

- **Absolute Air-Gap:** Nexus units are designed for deployment in secure enclaves or remote operational zones, allowing for a complete air-gap from the public internet.
- **Jurisdictional Control:** Because the infrastructure is mobile and independent, it can be sited precisely within the legal or physical jurisdiction required by the sovereign authority.
- **Operational Autonomy:** The customer owns the compute and the physical access, ensuring that model parameters and proprietary datasets remain within their chain of custody at all times.4. Infrastructure as a Foundation of Security

A sovereign compute stack is only as secure as the infrastructure that powers it. The most common point of failure for sovereign or "edge-secure" deployments is the reliance on the local utility grid. A fragile, public-facing utility tie creates a vulnerability; if the grid goes down, the secure compute goes dark.

Flux Core’s sovereign deployments are **power-agnostic**. We don't need a public utility hookup to operate. By integrating on-site microgrids—whether powered by high-reliability natural gas generators, hydrogen fuel cells, or solar/battery storage—we decouple the sovereign enclave from the public infrastructure grid.

Furthermore, our **zero-water liquid cooling** allows us to deploy these secure units in austere or remote locations where water-dependent cooling towers would be impossible or would necessitate external maintenance access, further reducing the physical footprint of the site and minimizing the number of external personnel required for maintenance.5. Conclusion: Owning the Compute, Data, and Power

Sovereignty is not achieved through software policy; it is achieved through architectural control. By merging high-density GPU compute with power-independent generation and zero-water thermal management, Flux Core provides the only viable model for truly sovereign AI.

We are currently scoping sovereign compute deployments for defense and enterprise partners. If your organization requires a secure, air-gapped, or geographically restricted compute environment, contact the Flux Core strategy team to review your site requirements.

White Paper: The Powered Land Bubble – Why the Real Estate Model of AI Infrastructure is Destined to Fail

Flux Core Data Systems

1. Executive Summary: The False Economy of "Powered Land"

The AI infrastructure market is currently experiencing a dangerous misalignment between real estate speculation and the reality of compute deployment. Across the United States, we are seeing 4MW of "powered land" marketed for $3 million annually—excluding electricity costs. This equates to $750,000 per megawatt, per year, for nothing more than dirt and an optimistic interconnection promise. The compute is not built, the chips are not racked, and the tenant is effectively paying a massive premium for the privilege of waiting five years in a grid queue.

This white paper examines why the "Powered Land as a Service" model represents a commercial real estate bubble rather than a sustainable infrastructure strategy. For operators and enterprise clients, the path to profitability is not found in renting the right to wait; it is found in owning the power generation and decoupling compute from legacy utility constraints.

1. The Anatomy of the Bubble

The fundamental flaw in the current powered land model is the assumption of perpetual grid scarcity. Landlords are underwriting decade-long leases against an interconnection bottleneck that is already eroding. They are betting that developers will remain trapped by the grid for the foreseeable future.

This model creates three distinct failures:

- **The Interconnection Trap:** By anchoring deployments to the grid, developers become hostage to utility timelines. A 5-year wait for a substation upgrade is an eternity in AI development.
- **The Valuation Gap:** The $750,000/MW annual ground lease is an unsustainable tax on compute. As behind-the-meter generation becomes the industry default, these lease valuations will collapse.
- **The Stranded Asset Risk:** When compute can be generated behind the meter, the land—which was valued primarily for its proximity to a grid tie—loses its primary competitive advantage. The rent rolls behind these long-term contracts are built on a foundation of artificial scarcity.

1. The Shift: Behind-the-Meter as the New Standard

The industry is rapidly pivoting. Data from current market trends suggests that behind-the-meter capacity in the United States could exceed 40GW by 2028. This move is not accidental; it is an economic necessity.

When a developer can generate power on-site—utilizing stranded natural gas, solar paired with battery energy storage, or other hybrid microgrid solutions—the "powered land" premium evaporates. The moment the power is sourced behind the meter, the need for a premium-priced ground lease disappears.

Flux Core has positioned its infrastructure strategy on the other side of this trade. We do not rent the utility's bottleneck; we develop our own energy independence.

1. Flux Core’s Alternative: Infrastructure as a Product, Not a Lease

Flux Core rejects the "Powered Land as a Service" model. Our approach is defined by architectural agility rather than real estate speculation:

- **Power Agnostic Deployments:** Our Nexus units are designed to run on whatever energy source is most cost-effective and available, whether that is stranded gas, renewable solar, or an existing grid tie.
- **Zero-Lease Operations:** By eliminating the need for premium-priced "powered land," we remove the pass-through markup on electricity and the burden of ground leases.
- **Thermal Efficiency:** Our proprietary closed-loop liquid cooling system uses zero outside water, allowing us to deploy compute in regions that are traditionally considered "unbuildable" by standard operators—such as arid areas with low-cost energy—further distancing us from the expensive, grid-dependent land market.

1. Conclusion: Stop Renting a Power Problem

The market is currently suffering from a bubble of optimism. Landlords are charging for the right to wait, but the next wave of AI infrastructure will belong to those who can generate their own power and deploy on their own timeline.

As behind-the-meter generation scales, the value of conventional powered land will inevitably correct. Flux Core is helping our partners move away from this bubble by delivering deployable, power-agnostic, and thermally efficient infrastructure.

Stop renting a power problem. Talk to the Flux Core strategy team about owning your energy, monetizing your data, and bypassing the grid-dependent real estate trap.