1. Offtake Agreements

Ask any lender what makes an AI data center bankable, and the answer is a signed offtake. Contracted demand is what turns a power site into a financeable asset, and it is the piece most developers cannot lock down. Flux Core leads with it. We hold a Memorandum of Understanding covering 50 to 100 MW of compute demand, and we are opening additional offtake slots now.

Who takes the capacity? A deliberately broad base, so no single market controls the schedule: neoclouds reselling GPU hours, enterprises running private models, universities and national labs, hospitals, municipalities, and AI platforms scaling inference. Each partner contracts for capacity inside a containerized Nexus deployment that comes online in months, with closed-loop liquid cooling that consumes zero outside water.

Because Flux Core is power agnostic, that capacity can sit on stranded gas, renewable solar with storage, or grid power, so an offtake is never held hostage by an interconnection queue.

If you are sourcing committed AI compute, or want to anchor a site as an offtake partner, reach out to Flux Core to review open megawatts.

---

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

Picture a 1.3MW rack hall that is already power-fed, already liquid-plumbed, and already thermally proven, waiting for nothing but your servers. That is what Flux Core hands a tenant on day one. Most powered shells stop at a slab and a switchgear lineup; ours arrives ready to absorb the densest AI clusters on the market.

The cooling is the difference. A closed-loop liquid loop circulates the same fluid continuously and draws zero outside water, so there is no cooling tower, no make-up water permit, and no evaporative loss to engineer around. Racks run hot-chip dense at Tier III reliability without the water and air-handling overhead that caps conventional colocation.

The shell is also power agnostic. Drop it on stranded natural gas, on solar paired with battery storage, or on an existing grid tie, and Flux Core matches the power path to the site. That is how a fully provisioned 1.3MW environment reaches edge and constrained locations the big colos skip.

Scaling GPUs and tired of waiting on a shell to be finished? Ask Flux Core about a deploy-ready powered shell up to 1.3MW.

---

3. Power as a Service

Interconnection queues in the largest US markets now stretch past five years, and that wait, not silicon, is what keeps AI capacity offline. Flux Core removes the wait entirely. With Power as a Service, we finance, build, own, and operate the generation and distribution, and you simply buy clean, dispatchable power for your compute as an operating expense.

No capital outlay. No utility study limbo. No balance-sheet exposure to generators, switchgear, or batteries. Flux Core is power agnostic, so the system underneath you can be stranded natural gas, solar paired with battery storage, a hybrid microgrid, or a grid tie, sized to your load profile and the site.

The cooling side stays just as clean. Closed-loop liquid cooling with zero outside water keeps PUE low and keeps your facility off local water supply, which matters in the arid and permitting-sensitive regions where cheap power tends to live. You pay a predictable rate; Flux Core monetizes the credits and tax structures behind it.

Stuck in an interconnection queue? Ask Flux Core how Power as a Service brings your compute online now.

---

4. Oil and Gas Well Deployments

The World Bank estimates the gas flared worldwide each year could power entire countries. On a single pad, that same wasted molecule can instead run GPUs. Flux Core deploys containerized data centers at the wellhead and converts stranded and associated gas into AI compute revenue, on site, with no pipeline required.

For an operator, the math is direct. Gas with no takeaway capacity, gas that is currently flared, vented, or sold at a steep basis discount, becomes the fuel for on-site generation that powers high-value compute and token output. It is a new revenue line from an existing asset, and reduced flaring improves the emissions profile at the same time.

Water is usually the constraint in a basin, and our closed-loop liquid cooling answers it by using zero outside water. The unit ships on a truck, runs power agnostic on that wellhead gas today and on renewables or grid later, and relocates as the field's production curve shifts.

Have gas with nowhere to go? Flux Core puts it to work as compute. Let's scope a wellhead deployment on your acreage.

---

5. Zero Carbon Footprint

Sustainability claims usually live in a marketing deck. Flux Core engineers ours into the hardware. A genuinely low-carbon AI data center has to solve generation, emissions, and cooling together, and here is exactly how we do each one.

Generation comes first. Where conditions allow, Nexus units run on renewable solar paired with battery energy storage, displacing fossil power outright. Where gas is used, we prioritize stranded and associated gas that would otherwise be flared or vented, so the deployment reduces emissions that were already happening rather than adding new ones.

Then we close the gap. Avoided-flaring value and carbon credit mechanisms are paired against the remaining footprint to push net emissions toward zero. Finally, cooling: a closed-loop liquid system that uses zero outside water removes evaporative loss and the energy penalty of air cooling, which is what holds PUE low and sends more of every watt to the GPUs.

Renewable-first generation, flare remediation, carbon credits, and water-free cooling, stacked. That is a credible path to zero carbon.

Need infrastructure that survives an ESG audit? Ask Flux Core to walk through the carbon stack.

---

6. Sovereign AI, Sovereign Data

A regulated bank, a defense command, and a national ministry share one requirement: the data and the model cannot leave their control. Public hyperscale cannot meet that, and that is precisely the gap Flux Core fills. We place dedicated GPU compute on the customer's own ground, inside a perimeter they own end to end.

Each Nexus unit is self-contained and deployable, which means sovereign compute can stand up in a secure facility, a research enclave, or a forward austere site without backhauling sensitive workloads to someone else's cloud. Residency, security classification, and air-gap requirements are satisfied by design, not by policy promise.

Power independence reinforces 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 not tethered to a fragile utility connection. Closed-loop liquid cooling with zero outside water lets dense clusters operate even where water and infrastructure are scarce.

Sovereignty means owning the compute, the data, and the power together. Talk to Flux Core about standing up sovereign AI on your terms.

---

7. College and University Data

A research team wins a grant to train a large model, then discovers the campus substation cannot deliver the power, and a utility upgrade is years out. That stall is happening at universities everywhere. Flux Core gives higher education a way around it, delivering dedicated research compute to campus without waiting on a capital expansion.

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

It also fits how universities actually operate. Power agnostic means the cluster can draw on existing grid capacity, on solar with battery storage, or on-site generation, aligning with both tight budgets and campus sustainability pledges. Closed-loop liquid cooling with zero outside water keeps utility costs and environmental reporting clean.

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

---

8. Powered Land Lease Rates — The Real Bubble

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.