

1. Offtake Agreements

Hyperscale growth is throttled by one constraint: not enough compute reaching the buyers who need it. Flux Core closes that gap with secured offtake agreements that turn megawatts into committed, monetizable demand. We currently hold a Memorandum of Understanding for 50 to 100 MW of compute demand, and we are actively expanding our offtake partner base across neoclouds, enterprises, universities, hospitals, municipalities, national labs, and AI platforms.

An offtake agreement with Flux Core means guaranteed access to AI-ready GPU capacity inside containerized data centers that deploy in 6 to 9 months, not the 2 to 4 years a traditional build demands. Each Nexus R1000 unit delivers up to 1MW of compute, cooled by a closed-loop liquid cooling system that uses zero outside water, and powered by whatever the site offers. Our power-agnostic model draws on stranded gas, renewable solar paired with battery storage, or existing grid connections, so capacity is never stranded waiting on a substation.

For compute buyers, this converts an unpredictable supply market into contracted certainty. For Flux Core, every offtake commitment de-risks deployment and accelerates the next site.

Looking to secure AI compute capacity through an offtake agreement? Contact Flux Core Data Systems and Energy to discuss available megawatts.

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

Most powered shells hand you four walls and a power feed, then leave the hardest problems, cooling and density, to you. Flux Core delivers a fully provisioned bare metal environment engineered to support up to 1.3MW of IT load per deployment, ready for the highest-density AI and GPU workloads from day one.

Our bare metal provisioning pairs a powered, AI-ready shell with a closed-loop liquid cooling system that uses zero outside water, eliminating the water-permitting and evaporative-loss problems that stall conventional hyperscale builds. Tenants bring their own hardware or lease ours, and drop straight into a thermally validated, high-density environment rated to Tier III standards for 99.982% uptime.

The shell is power agnostic. Whether your site runs on stranded natural gas, renewable solar with battery storage, or an existing grid tie, Flux Core configures the power path to match. That flexibility is what lets a 1.3MW powered shell go live in remote, edge, or power-constrained locations where traditional colocation simply cannot reach.

The result is a turnkey, deployable foundation for AI infrastructure: dense, liquid-cooled, water-free, and powered by whatever the site demands.

Need a powered shell for high-density GPU deployment? Talk to Flux Core about bare metal provisioning up to 1.3MW.

3. Power as a Service

The bottleneck in hyperscale AI is no longer chips. It is power, and the multi-year grid interconnection queues standing between a project and live capacity. Flux Core solves this with Power as a Service: we own, build, and operate the energy system, and you buy reliable power for your compute without the capital burden or the wait.

Flux Core is power agnostic by design. We generate and deliver electricity from stranded natural gas, renewable solar paired with battery storage, hybrid microgrids, or existing grid connections, matching the source to whatever the site demands. That flexibility frees AI deployments from the grid capacity that is already booked through 2030 in most primary markets.

Because our Nexus units use a closed-loop liquid cooling system with zero outside water, the energy footprint stays efficient and predictable, with low PUE and no dependence on local water resources. You contract for power and cooling as an operating expense, while Flux Core captures the energy credits and tax structures that make the underlying system economical.

Power as a Service turns the single hardest problem in AI infrastructure into a simple monthly line item.

Want compute that never waits on the grid? Ask Flux Core about Power as a Service.

4. Oil and Gas Well Deployments

Every flared and stranded gas well is wasted energy and a wasted balance sheet. Flux Core converts that liability into AI compute revenue by deploying containerized data centers directly at the wellhead, powered by the gas that would otherwise be burned off or sold at a discount.

Our power-agnostic Nexus units run on stranded and associated natural gas with on-site generation, transforming gas with no pipeline takeaway into high-value GPU compute and tokens. Operators monetize molecules that currently have no market, reduce flaring, and create a new revenue stream from existing assets, no new drilling required.

Each deployment uses a closed-loop liquid cooling system that uses zero outside water, which is decisive in arid oilfield basins where water is scarce, expensive, and heavily regulated. The containerized form factor means a unit can be trucked to a remote pad and brought online in 6 to 9 months, then relocated as production profiles change.

For oil and gas operators, this is energy monetization without midstream dependency. For AI buyers, it is reliable, low-cost compute capacity in places the grid will never reach.

Sitting on stranded or flared gas? Flux Core turns it into AI compute. Contact us about wellhead deployments.

5. Zero Carbon Footprint

A data center can be both high-density and low-carbon, but only if energy and cooling are engineered together. Flux Core achieves a zero carbon footprint through a deliberate stack of renewable power, captured emissions, and water-free thermal design.

It starts with our power-agnostic architecture. Where the site allows, Flux Core powers Nexus units with renewable solar paired with battery energy storage, displacing fossil generation entirely. Where natural gas is used, including stranded and flared gas that would otherwise be vented, we capture the value of avoided flaring and pair generation with emissions offsets and carbon credit mechanisms to drive net emissions toward zero.

Cooling is the second half of the equation. Our closed-loop liquid cooling system uses zero outside water, eliminating evaporative water consumption and the energy penalty of traditional air cooling. The result is a low PUE that means more of every watt goes to compute rather than to overhead.

Layered together, renewable generation, stranded gas remediation, carbon credits, and water-free closed-loop cooling produce AI infrastructure that meets aggressive sustainability and ESG targets without sacrificing density or uptime.

Building toward net zero AI? Ask Flux Core how we engineer a zero carbon footprint.

6. Sovereign AI, Sovereign Data

When data cannot leave the jurisdiction, the building, or the theater, hyperscale cloud is not an option. Flux Core delivers sovereign AI and sovereign data through containerized data centers that put dedicated GPU compute exactly where the mission and the law require it.

Our Nexus units are deployable, self-contained AI infrastructure that governments, defense organizations, and regulated enterprises can site on their own ground, fully under their own control. Workloads and data stay inside the customer's perimeter, supporting strict residency, security, and air-gapped requirements that public cloud cannot satisfy.

Because each unit is power agnostic, sovereign compute can stand up in remote bases, secure facilities, or austere edge environments using stranded gas, renewable solar with battery storage, or local grid power, whatever the site demands. Independence from the grid is independence for the mission.

The closed-loop liquid cooling system uses zero outside water, so even water-constrained or sensitive sites can host dense GPU clusters. Tier III design delivers 99.982% uptime, and deployment takes 6 to 9 months rather than years.

Sovereign AI means owning your compute, your data, and your power. Talk to Flux Core about sovereign, deployable AI infrastructure.

7. College and University Data

Universities are sitting on world-class AI research ambitions and grid connections that cannot deliver the power those ambitions require. Flux Core gives higher education a faster path to dedicated research compute without waiting years for a campus expansion or a utility upgrade.

Our containerized Nexus data centers bring up to 1MW of GPU capacity to campus for research computing, high-performance computing, and AI model training, deployable in 6 to 9 months. Research directors get sovereign control over sensitive datasets, grant-funded workloads, and student access, with compute that stays on institutional ground rather than in a shared public cloud.

Flux Core is power agnostic, so a campus can power its cluster from existing grid capacity, renewable solar paired with battery storage, or on-site generation, fitting both the budget and the sustainability commitments that universities increasingly carry. The closed-loop liquid cooling system uses zero outside water, an important advantage for institutions managing utility costs and environmental reporting.

Whether the goal is a flagship AI initiative, a shared HPC resource, or secure research data infrastructure, Flux Core delivers the density and uptime academic workloads demand.

Expanding research computing on campus? Contact Flux Core about university AI infrastructure.