



OFFERINGS BROCHURE · 1 MW+ CONTAINERIZED DATA CENTERS

Three Paths to 1 MW+ Compute

Powered Shell · Bare Metal · Turnkey

Modular 1 MW+ AI infrastructure, engineered for accelerated deployment.
Client-owned GPUs, turnkey compute, or full lease-and-operate — you choose.

\$10-12M / MW

POWERED SHELL

Client brings GPUs

\$50-75M / MW

BARE METAL

B200 / B300 / GB300

\$12-20M+ / yr

TURNKEY

Lease or own • FC operates

Three Paths to 1 MW+

Choose the depth of Flux Core engagement that fits your capital and control preferences

THE THREE OFFERINGS AT A GLANCE

Flux Core delivers modular 1 MW-and-larger AI data centers as containerized, closed-loop-cooled infrastructure built for rapid deployment. Whether you already own the accelerators, want the full compute stack, or prefer to lease and let us operate the asset, one of the three paths below aligns with how you want to deploy 1 MW+ of compute.

01

POWERED SHELL

You bring GPUs

\$10-12M / MW

Best for: Hyperscalers, GPU cloud operators, and enterprises with existing GPU procurement.

INCLUDES

- Container envelope
- Generators & switchgear
- Closed-loop liquid cooling
- Zero on-site water
- Power distribution
- Turnkey commissioning

GPU hardware · network fabric · software

02

BARE METAL

Client-ready compute

\$50-75M / MW

Best for: AI developers, GPU cloud operators, and enterprises that want compute-ready infrastructure.

INCLUDES

- Everything in Powered Shell
- NVIDIA B200 / B300 / GB300
- Racks & rack integration
- Power to the die
- Cooling to the accelerator
- Bring-up & burn-in

AMD & alternative accelerators quoted separately

03

TURNKEY

We operate it

\$12-20M+ / yr

Best for: Investors, hospitals, universities, and municipalities that want compute revenue without operating risk.

INCLUDES

- Lease or own the asset
- FC operates and manages
- Compute offtake included
- Revenue-share options
- 24/7 NOC & SOC
- Long-term O&M

Revenue depends on offtake mix and market rates

HOW TO CHOOSE

Powered Shell if you already have GPUs and want the fastest, cheapest path to a site-ready deployment. **Bare Metal** if you want a full compute stack that boots on “day one” without procuring silicon yourself. **Turnkey** if you want to own or lease the asset and let Flux Core drive the revenue engine.

1 MW+

CAPACITY

Scalable

\$10–12M / MW

PRICE

Turnkey install

Zero

ON-SITE WATER

Closed-loop

90–120 DAYS

TIME-TO-POWER

Site-ready to live

The **Powered Shell** is the fastest path to a deployable 1 MW+ AI site. Flux Core builds the containerized envelope, integrates the entire power and cooling stack, and delivers a GPU-ready site — you install and own the accelerators. Ideal for customers with existing GPU procurement, tight timelines, and preference for maximum control over their compute stack.

WHAT'S INCLUDED

CONTAINER ENVELOPE

Weather-sealed, seismic-rated ISO-form containers engineered for 1 MW+ per unit. Modular by design — add more.

ON-SITE POWER GENERATION

Natural-gas or diesel generator sets sized to full data center load, plus paralleling switchgear and load banks.

CLOSED-LOOP LIQUID COOLING

Rear-door heat exchangers with a fully recirculating coolant loop. **Zero on-site water consumption** — dry cooler rejects heat to atmosphere.

POWER DISTRIBUTION

480V three-phase distribution, 3000A busway, redundant PDUs to the rack. Ready for air- or liquid-cooled deployments.

TURNKEY COMMISSIONING

Complete site build, utility interconnect, factory acceptance, and on-site commissioning. Ready for your GPUs on delivery.

CLIENT-SUPPLIED / EXCLUDED

CLIENT PROVIDES GPUS

The customer is responsible for procuring and installing NVIDIA, AMD, or custom accelerator hardware.

CLIENT PROVIDES NETWORK FABRIC

Switch fabric, cross-connects, and network operations are outside the powered shell scope.

CLIENT PROVIDES SOFTWARE

Orchestration, workload management, tenancy, and customer-facing platforms are outside scope.

LOCATION ADJUSTMENTS

Utility tap fees, permitting, land, and site prep are passed through at cost — quoted per site.

OPTIONAL ADD-ONS

BESS, solar hosting, extended warranty, and managed services available as separate scope items.

WHY POWERED SHELL

Fastest path to deployment for GPU-holders. Predictable installed cost. Modular by design.

Zero water consumption removes the largest permitting friction point in AI data center siting.

\$50–75M / MW

PRICE RANGE
NVIDIA config

256–512

B200 / B300
Per 1 MW block

576

GB300 GPUs
Per 1 MW container

Ready to Run

DAY ONE
Bring workloads only

The **Bare Metal** tier delivers the full compute stack: everything in the Powered Shell plus factory-integrated NVIDIA accelerators — racks, cabling, cooling to the die, and burn-in — ready to run workloads on delivery. Configurations scale from dense inference to “day one” training on NVL72 rack-scale platforms.

NVIDIA ACCELERATOR CONFIGURATIONS · PER 1 MW BLOCK

INDICATIVE

PLATFORM	GPUs / 1 MW	DENSITY	TARGET WORKLOAD	PRICE / MW
NVIDIA B200	256 units	60 kW/rack	Training + inference	\$50–58M
NVIDIA B300	512 units	Blackwell Ultra	High-density training	\$70–74M
GB300 NVL72	576 GPUs	130 kW/rack liquid	Frontier training	\$68–75M

AMD & OTHER ACCELERATORS

AMD MI355X, Instinct-family, and custom/ASIC deployments quoted separately.
Our shell supports 40–132+ kW/rack liquid — forward-compatible with next-gen silicon.

WHAT BARE METAL ADDS ON TOP OF POWERED SHELL

NVIDIA silicon integrated at the factory — baseboards, DGX or HGX chassis, NVL72 rack-scale platforms. **Rack integration & structured cabling.** **Cooling to the die** with rear-door and direct-liquid options. **Full bring-up, burn-in, and diagnostics** so the cluster passes acceptance testing before it leaves the yard. Network fabric and customer software remain client-supplied or quoted as an add-on.

Turnkey · Lease or Own

Full asset. Flux Core operates and manages. Revenue engine included.

\$12-20M+

PER 1 MW / YR

Compute revenue

Lease

OR OWN

Client elects

FC Operates

24/7 NOC

Turnkey management

Long-Term

O&M

Enterprise-grade SLAs

The **Turnkey** tier is the full revenue engine. Flux Core builds, populates, and operates a 1 MW+ AI data center on behalf of an investor, enterprise, hospital, university, or municipality. The client elects **Lease** (predictable annual cost, FC owns) or **Own** (client owns the asset, FC operates under a long-term O&M agreement). Compute is monetized through offtake agreements that Flux Core structures on the client's behalf.

OPTION A

LEASE THE DATA CENTER

Client pays annual lease. FC owns and operates.

Structure	Annual lease + revenue share
Ownership	Flux Core
Operator	Flux Core – 24/7 NOC
Client Capex	Initial engineering fee
Client Upside	Compute offtake resale rights
Best For	Corporate offtakers · landlords

OPTION B

OWN THE DATA CENTER

Client acquires asset. FC operates under O&M.

Structure	Capital purchase + O&M contract
Ownership	Client
Operator	Flux Core – turnkey
Client Capex	\$50-75M / MW
Client Upside	Full compute revenue – O&M
Best For	Investors · hospitals · universities

REVENUE ENGINE · \$12-20M+ PER 1 MW PER YEAR

At current market rates for AI compute, a fully utilized 1 MW block conservatively generates **\$18-24M+ per year** in gross compute revenue. Depending on the chosen structure (lease vs own), operating costs, and market pricing, net revenue to the asset holder typically ranges from **\$12-20M+ per year per MW**. Larger deployments (2 MW, 5 MW, 10 MW+) scale revenue proportionally and improve unit economics.

Let's Build Together

All figures are indicative — every deployment is scoped to the site

ESTIMATE DISCLAIMER · PLEASE READ

All pricing, capacity, revenue, and timing figures presented in this brochure are estimates only. Final quotes are prepared on a per-site basis after a scoping engagement and may vary materially from the ranges shown. Factors that can affect any given deployment include, but are not limited to:

- **LOCATION & SITE CONDITIONS**
Utility availability, tap fees, permitting, soil, seismic, environmental.
- **MARKET CONDITIONS**
GPU pricing, commodity indices (steel, copper), lead times, freight.
- **EXISTING INFRASTRUCTURE**
Available power, water access, network, pad readiness, existing electrical service.
- **DEPLOYMENT SCALE**
Unit economics improve with scale — multi-MW deployments carry different pricing.
- **CHOSEN CONFIGURATION**
GPU platform, cooling density, redundancy tier, financing structure.
- **REGULATORY & INTERCONNECTION**
Utility interconnect timing, tariffs, and regulatory approvals.

Ready to Scope Your Deployment?

A 30-minute scoping call is enough for Flux Core to produce a site-specific quote, financing options, and deployment timeline for your project.

DYLAN CHANG

Co-Founder & Chief Strategy Officer
dylan@fluxcore.com

FLUX CORE DATA SYSTEMS AND ENERGY

Fresno, California · Deployments nationwide
www.fluxcore.energy | Rev 1.0 · July 2026