

GRIDFACTORY

Own the machines that power AI.

Green-energy H100 + B200 infrastructure

Compute-as-a-Service · Modular parks · Tokenised and lease-to-own participation

01 / THE OPPORTUNITY

Power-connected compute is the scarce asset.

AI demand expands faster than grid connections, permitting and high-density infrastructure. GridFactory moves training-class compute directly into renewable energy parks.

POWER MOAT

PARK → GPU

Deploy where land, interconnection and renewable generation already exist.

PRODUCT

32 + 32

A standardized 64-GPU module combining compute, 400G fabric, cooling, storage and software.

REVENUE

CaaS

Reserved training capacity, enterprise leases and sovereign AI contracts—not low-value desktop-GPU resale.

THE INVESTOR PROMISE

Enter the physical AI layer before power-connected compute becomes institutional-only.

Long-horizon thesis

Personal, industrial and defence robotics may expand the AI-enabled economy far beyond today's semiconductor valuations. GridFactory treats this as a strategic vision—not as present TAM, contracted revenue or guaranteed value.

02 / WHAT EXISTS TODAY

Separate the asset, the plan and the ambition.

Classification	Current statement	Investor evidence required
Operator-provided	Two Romanian park opportunities secured	LOI/PPA/site-control documents
Designed	64-GPU mixed module and 400G fabric	Final engineering and supplier BOM
Forecast	H100/B200 rates and utilization	Customer LOIs/reserved contracts
Planned	Tokenisation and lease-to-own	UAE/EU legal opinions and structure
Aspirational	UAE park and multi-park network	Pipeline and execution milestones

RULE FOR INVESTOR MATERIALS

Do not describe Container 01, tokenisation, payouts, Tier IV certification, 40 MW availability, PUE 1.18 or zero-carbon operation as live unless the data room contains current documentary evidence.

03 / THE MODULE

One container. Two premium GPU fabrics.

RACKS 1-4 · H100 SXM

32 H100

32 GPUs across four eight-GPU nodes. Mature training ecosystem, strong availability and reserved-capacity pricing reference.

RACKS 5-8 · B200 SXM

32 B200

32 GPUs across four eight-GPU nodes. Higher compute density, higher power envelope and premium training-rate potential.

Layer	H100 fabric	B200 fabric
Reference throughput	989 TFLOPS BF16 dense/GPU	4.5 PFLOPS FP8 dense/GPU
Allocated system power	1.28 kW/GPU	1.79 kW/GPU
Commercial base rate	\$3.90/GPU-hour	\$8.50/GPU-hour
Primary workload	Reserved LLM training	Premium training/inference

Precision formats are not directly comparable. Published peak throughput is not achieved application throughput; customer benchmarks govern commercial performance.

04 / THE FORMULA

Every return begins with five auditable inputs.

Line	Formula	Base assumption
H100 revenue	GPUs x rate x 720h x utilization	32 x \$3.90 x 720 x 85%
B200 revenue	GPUs x rate x 720h x utilization	32 x \$8.50 x 720 x 88%
Energy	IT kW x PUE x 720h x \$/kWh	1.25 PUE x \$0.11/kWh
Operating cost	8% revenue + fixed module OpEx	\$6,500/month/module
Module EBITDA	Revenue - energy - operating cost	Before finance, tax and D&A
Simple payback	Installed CapEx / monthly EBITDA	No terminal value included

INSTALLED CAPEX FORMULA

H100 count x \$39,675 + B200 count x \$65,000 + \$469,500 shared module infrastructure. Planning assumptions remain subject to supplier quotations and final engineering.

05 / SENSITIVITY

Profit is governed by rate, utilization and power.

Scenario	Gross / month	EBITDA / month	CapEx	Simple payback
Downside	\$143,770	\$110,914	\$3,819,100	34.4 mo
Base	\$248,717	\$212,594	\$3,819,100	18.0 mo
Upside	\$298,898	\$261,809	\$3,819,100	14.6 mo

BASE UTILIZATION

85–88%

H100 85% · B200 88% · reserved-training weighted mix.

BASE ENERGY

\$0.11

\$0.11/kWh · PUE 1.25 · system-level power allocation.

Model boundary

EBITDA here deducts modeled energy, an 8% operating reserve and \$6,500/month module overhead. It excludes financing, tax, depreciation, legal/token costs and customer-acquisition ramp. Final numbers must be rebuilt from supplier quotes, PPA terms and signed demand.

06 / SCALE PLAN

Add modules only after demand unlocks capital.

Year	Modules	GPUs	Revenue	EBITDA	Installed CapEx
Y1	1	64	\$2.5M	\$2.1M	\$3.8M
Y2	2	128	\$5.5M	\$4.7M	\$7.6M
Y3	4	256	\$11.9M	\$10.1M	\$15.3M
Y4	6	384	\$18.2M	\$15.6M	\$22.9M
Y5	9	576	\$27.9M	\$24.0M	\$34.4M

Y5 PLANNING CASE

576 GPUs

Nine mixed modules, 576 training-class GPUs and 90% blended utilization target.

EXPANSION GATE

GATED

No next module before power readiness, supplier delivery and reserved demand are evidenced.

This is a capacity scenario, not a forecast commitment. It repeats current mixed-module assumptions and excludes procurement inflation, financing, tax and working-capital drag.

07 / DEPLOY IN STAGES

Capital enters as hardware becomes productive.

Stage	Deployment	Indicative raise	Revenue trigger
A	Foundation + 16 H100	\$1.27M	First reserved H100 invoices
B	Add 16 H100	\$0.70M	Complete H100 fabric
C	Add 16 B200	\$1.14M	Premium mixed capacity
D	Add 16 B200	\$1.14M	64-GPU module complete
TOTAL	32 H100 + 32 B200	~\$4.25M	Full mixed module

WHY STAGED CAPITAL MATTERS

Each stage creates a diligence checkpoint: supplier delivery, power readiness, cluster acceptance, customer utilization and invoicing. Earlier infrastructure can support later deployment, but reinvestment is an option—not an automatic assumption.

08 / GO TO MARKET

Sell commitments, not cheap spot hours.

SOVEREIGN AI

Domestic/private training clusters with data-residency requirements.

AI LABS

Reserved multi-node training and fine-tuning capacity.

ENTERPRISE

Dedicated capacity, platform support and predictable billing.

AGENTIC + ROBOTICS

Persistent inference, simulation and autonomy workloads.

COMMERCIAL PRIORITY

1. Anchor reserved contract → 2. Cluster acceptance → 3. Expand inside the park → 4. Replicate across parks

The previous \$0.55–\$0.92/GPU-hour consumer resale model is removed. H100/B200 economics are modeled from premium infrastructure rates and must be validated through customer commitments.

09 / PARTICIPATION

The same machine can support several capital instruments.

EQUITY

Platform and multi-park enterprise-value upside.

PROJECT DEBT

Coupon, maturity and defined security package.

REVENUE SHARE

Contractual share of distributable CaaS revenue.

INFRASTRUCTURE-BACKED

Capital linked to named hardware or a module.

TOKENISED GRID UNITS

Digital representation subject to regulatory structuring.

OWN-TO-LEASE

Investor-owned equipment leased into operations.

No structure should be marketed as guaranteed, fixed or risk-free until legal documents, collateral controls, waterfall and jurisdictional approvals are complete.

10 / THE SECOND RETURN ENGINE

Cash yield proves the asset. Software can reprice the company.

ASSET VALUE

HARDWARE

GPUs, storage, networking, container and power equipment support recoverable infrastructure value.

OPERATING VALUE

REVENUE

Reserved contracts and utilization create predictable CaaS cash flow.

PLATFORM PREMIUM

MULTIPLE

Software, customer access, energy-park pipeline and tokenisation may support an ARR multiple.

Interactive exit model

The portal allows investors to select 2×–12× projected ARR, reinvestment, dilution and exit timing. The default 7× is an illustrative midpoint—not a current GridFactory valuation. Pre-revenue companies should normally carry a material discount to scaled public AI-cloud comparables.

Comparable signal

CoreWeave reported \$5.13B of 2025 revenue and 98% of revenue from committed customer arrangements. Its market value can inform—but never determine—GridFactory's eventual multiple. Source: CoreWeave 2025 Form 10-K; market value changes daily.

11 / VALUE BRIDGE

Operating cash flow plus a software-enabled exit.

Modeled dilution	4x ARR	7x ARR	10x ARR
15%	25.4x	43.4x	61.3x
25%	22.6x	38.4x	54.3x
40%	18.4x	31.0x	43.7x

INFRASTRUCTURE DISTRIBUTION

\$301K/YR

Illustrative 45% of pro-rata base EBITDA:
\$300,596/year on a \$1.0M participation.

PLATFORM EXIT BASIS

Y5 ARR

Y5 planning ARR of \$26.9M across nine modules,
adjusted for selected multiple and dilution.

Method note

Outputs combine five years of flat illustrative distributions with a pro-rata platform exit. They exclude financing, taxes, preferences, fees, further capital and timing differences. These are scenarios, not an offer, valuation, forecast or guaranteed liquidity.

12 / EXECUTION

Prove one module. Then earn the right to scale.

Gate	Milestone	Evidence
1	Park and power diligence	Site control, PPA, interconnection
2	Stage A procurement	Quotes, escrow, delivery schedule
3	Cluster acceptance	Burn-in, 400G tests, performance
4	Commercial proof	Signed capacity and first invoices
5	Campus expansion	Utilization and reinvestment decision

Key risks

- + GPU procurement, pricing and architecture obsolescence.
- + Customer concentration and utilization below plan.
- + Power-price, curtailment, cooling and availability risk.
- + Tokenisation, securities, MiCA and UAE regulatory risk.
- + Execution risk across two countries and multiple suppliers.

IMMEDIATE INVESTOR ASK

Fund and diligence Stage A: foundation infrastructure plus the first 16-GPU H100 cluster. Full mixed-module ambition is approximately \$4.25M, released against documented deployment gates.