MICROLINK DATA CENTERS

The Heat Return Pod

A ring fenced SPV that owns, builds, funds and operates 2 MW compute pods behind the meter at industrial host partners. Contracted take or pay compute revenue; recovered heat as the siting moat.

2 MW	87	174 MW	29.4%
POD UNIT	PODS BY YEAR FIVE	PROGRAMME	LEVERED IRR

What is in this pack

The constraint is powered capacity

AI compute demand outstrips deliverable capacity. The binding constraint is energised, liquid cooled, connected megawatts near where they are needed.

Greenfield arrives in years.

Interconnection queues, permitting and a construction draw put new capacity years out. The market prices that scarcity in lead time.

The pod arrives in months.

Factory build, drop, three connections, commission. Under 20 weeks from a signed partner to live once the design is locked.

Demand already exceeds supply.

Compute is contracted through an aggregation partner whose committed take or pay offtake exceeds our deployment capacity. Pods deploy against signed demand, never speculatively.

Years	Months
GREENFIELD PATH TO FIRST COMPUTE	HEAT RETURN POD PATH
Over demand	**Monthly**
CONTRACTED OFFTAKE CHANNEL	NEW SITES AT RUN RATE

About 1,500 sites can take a pod today

Roughly 700 district energy systems and 16,000 treatment plants in the country. The grid connection, the land and the site services are already standing.

~700

DISTRICT ENERGY SYSTEMS

~16,000

TREATMENT PLANTS

~1,500

SITES QUALIFY TODAY

~1,250

DIGESTERS HELD AT 37°C (99°F)

THE THREE SITE TESTS

Headroom	~2 MW free continuously behind the meter
Heat sink	Year round hot water loop, digester at 37°C (99°F), or a small heat pump lift
Fibre	Carrier or research drop on a maintainable path

Heat opens the gate. The pod returns ~1.6 MW of recovered heat to the partner network, 85% of delivered volume free. That is why power is agreed behind the meter, and why the sector associations, IDEA, WEF, NACWA and the American Biogas Council, co publish our interconnection specs. Hosts are partners, never customers.

One identical unit serves both channels.

Contracted before it ships

Take or pay against a committed kW floor. Power is metered and passed through at cost with no markup, so the SPV carries no power price exposure.

$154
PER KW PER MONTH BASE RENT

95%
CONTRACTED AVAILABILITY

3%
FIXED ESCALATION FROM YEAR TWO

At cost
POWER PASS THROUGH

Compute capacity is contracted through an aggregation partner whose committed offtake demand exceeds our deployment capacity. The base rent sits inside the market comp band of $148 to $156 per kW per month for liquid cooled high density capacity.

REVENUE PER POD, YEAR ONE

Compute base rent	$3.20M, 1,820 kW at $154 and 95%
Heat revenue	$0.09M, 15% of delivered volume at avoided boiler cost
Total revenue	$3.28M
Tenant silicon	13 racks at ~140 kW, GB300 NVL72 class, tenant furnished

Heat is the moat. ERE is the number.

A good PUE says you waste little running the building. It says nothing about the heat. ERE below 1.0 means a site returns more usable energy than it loses.

<0.5

ERE TARGET, HEAT RECOVERED

1.12

DESIGN PUE, 1.09 TRI GENERATION

~1.6 MW

THERMAL TO THE PARTNER LOOP

85% free

REMAINDER AT AVOIDED BOILER COST

Why it is a moat.

Heat return is the reason ~1,500 fence lines open, the reason power is agreed behind the meter, and the reason the channel exists at all.

Whoever writes the spec owns the category.

Rollout runs through IDEA, WEF, NACWA and the American Biogas Council: co published interconnection specs, pilot cohorts, a joint waitlist.

Reported together, always.

ERE alongside PUE, from EN 1434 class metering at the host partner boundary. Heat is delivered below 45°C (113°F), digester grade and district grade with a small lift.

Built, tested and cabled before it leaves the floor

Two 13.716 m (45 ft) HC containers, factory built and witness tested by Vertiv. On site it takes three connections: power, coolant, fibre.

UNIT SUMMARY

Pod capacity	2 MW behind the meter
Compute	~13 GB300 NVL72 racks, ~936 Blackwell Ultra GPUs, tenant furnished
Heat returned	~1.6 MW to the partner loop, below 45°C (113°F)
Cooling	Direct to chip liquid, dry cooler always on the rejection path
Efficiency	ERE reported alongside PUE
Signed partner to live	Under 20 weeks once the design is locked

Trade work,

not construction. It arrives on a truck, is craned onto its pad, and connects at three points.

Hardware,

not real estate. The pod scales like a manufactured product, and the factory slot sets the schedule.

Three loops. One of them is always open.

Three loops carry heat from the chip to where it is needed. The third always has somewhere to go, so compute never waits and the partner is never the bottleneck.

Loop 01, off the silicon.

Direct to chip cold plates to a closed, treated secondary loop, at the density modern racks demand.

Loop 02, across the site.

The CDU moves that heat through to the facility heat exchangers, clean and contained, N+1 throughout.

Loop 03, to the partner.

Heat goes to the partner who needs it, metered EN 1434 at the boundary. When demand drops, the dry cooler rejects it instead.

DRY COOLER FIRST

The rejection plant is sized to carry the full 2 MW load on its own, so availability never depends on the heat being taken. Utilization never depends on host demand: compute revenue is decoupled from the thermal side entirely. That is the underwriting claim.

One boundary, one meter.

The partner side is a single heat exchanger tie in with an EN 1434 class meter. Everything else lives inside the unit.

87 pods by year five

Factory build is the throttle. New pods contribute a half year in their delivery year; each lands in under 20 weeks once the design is locked.

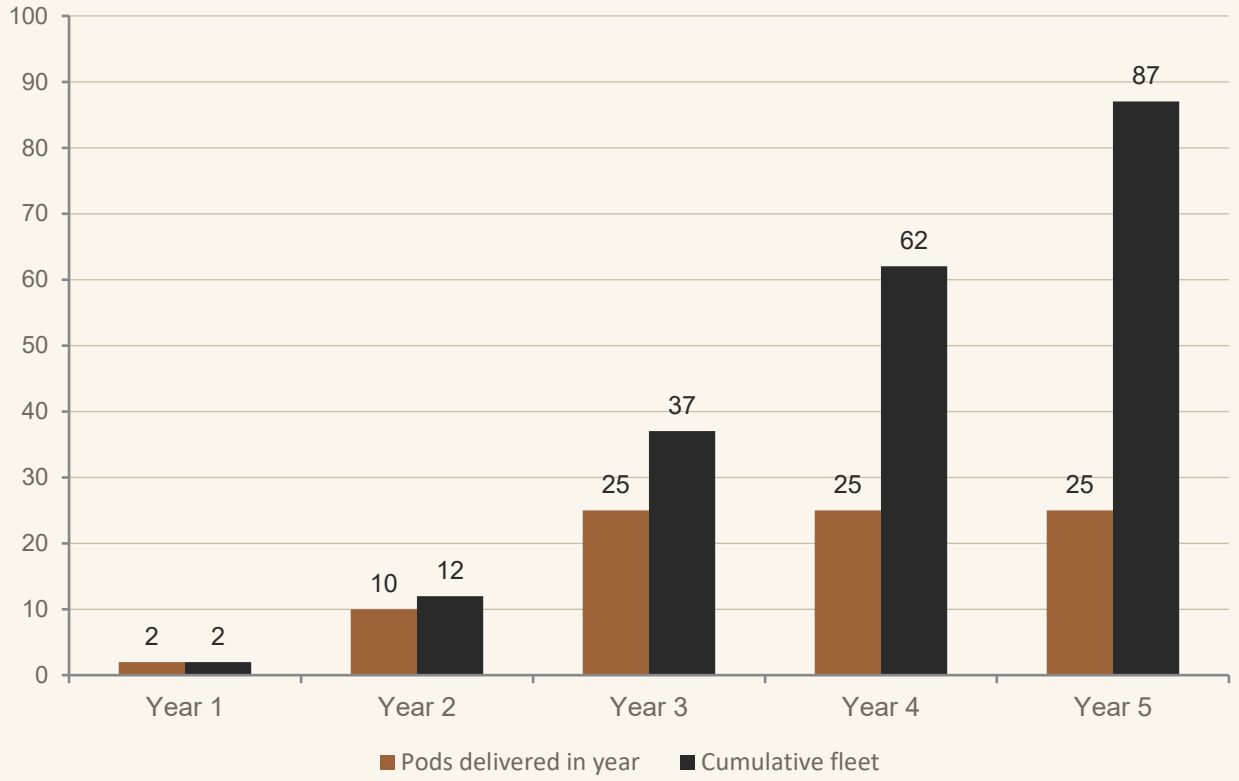

PROGRAMME

Deliveries	2, 10, 25, 25, 25 pods per year
Year five fleet	87 pods, 174 MW
Run rate	~25 pods per year from year three
Site runway	~1,500 qualifying today

A $21M asset at a 69% margin

Underwritten at $10.5M per MW, the mid of the $8.5M to $12.5M study band. Power passes through at cost, so revenue is nearly pure rent against a fixed opex base.

PER POD, YEAR ONE

Capex, installed	$21.0M at $10.5M per MW
Compute base rent	$3.20M
Heat revenue	$0.09M at avoided boiler cost
Opex, all in	$1.01M
EBITDA	$2.27M, 69% margin
Life	20 year depreciation, $400K CDU and pump overhaul in year eight

$2.27M

EBITDA PER POD, YEAR ONE

69%

EBITDA MARGIN

Opex is the NOC allocation, regional hands, maintenance, insurance, the site licence to the host partner, connectivity, tax and G&A. The NOC share falls per pod as the fleet grows. Revenue and opex both escalate at 3% from year two, so the margin holds.

17.3% unlevered. 29.4% levered.

Ten year hold, year ten exit at a 6.25 cap on EBITDA, pre tax basis shown. Debt is 75% at 7.5%, 15 year fully amortising.

17.3%
UNLEVERED IRR, TEN YEAR

29.4%
LEVERED IRR AT 75/25

7.6 yrs
LEVERED PAYBACK PER POD

$5.25M
EQUITY PER POD

DSCR builds from 1.27 in year one to 1.66 by year ten. Year one sits below the 1.35 covenant and is addressed in structuring.

SENSITIVITY, UNLEVERED YIELD ON COST

Capex $M/MW	$135	$154	$165	$185
8.5	11.0%	13.4%	14.7%	17.1%
10.5 base	8.9%	**10.8%**	11.9%	13.9%
12.0	7.8%	9.5%	10.4%	12.1%
12.5	7.5%	9.1%	10.0%	11.7%

Base rent columns in $ per kW per month. Base case marked.

Ring fenced SPV. MicroLink does everything.

The SPV owns the pods and the site agreements. Host partners take no ownership and no operating role; the compute tenant furnishes the silicon.

WHO DOES WHAT

MicroLink	Owner, builder, capital partner, deployer and operator through the ring fenced SPV
Vertiv	Factory integrated pod: build, fit out, witness test, ship
NVIDIA	Rack reference designs and the allocation path for the tenant fleet
Host partners	Power behind the meter at the agreed rate, loop tie in, heat offtake. Partners, never customers
Offtake channel	Aggregated compute demand, take or pay against a committed kW floor

THE ASK

Fund the first cohort

This raise funds the first cohort of a staged programme: the pods reaching first compute in years one and two, deployed against contracted demand at partner sites already in the waitlist. Each cohort stands on its own returns. Terms, structure and diligence materials available in conversation.

12	24 MW	29.4%	<20 wks
PODS, FIRST COHORT	DEPLOYED	LEVERED IRR, BASE	SIGNED TO LIVE