An extra 2 MW of compute that heats the town

A repeatable pod dropped behind a district energy site. It computes, and returns the heat to the network. MicroLink owns it, builds it, funds it and operates it. With NVIDIA, Vertiv, Hydra Host and IDEA.

~1.6 MW	$0	100 kW	<20 wks
HEAT TO YOUR NETWORK	CAPITAL FROM YOU	COMPUTE IF WANTED	SIGNED TO LIVE

What is in this pack

The most valuable heater ever built

It just happens to compute first. A chip turns almost 100% of the power it draws into heat. Every data center today throws that heat away. We treat it as a second product.

100%

ALMOST EVERY WATT LEAVES THE SILICON AS HEAT

A boiler has one output.

Fuel in, heat out. One product, and you buy the fuel for it.

A chip gives you two from the same watt.

The compute the tenant pays for, and the heat you can use. Free heat, on power you already sold at the agreed rate.

The industry is moving to liquid. That is the opening.

Racks are outgrowing what air can cool. The industry made this transition for its own reasons; your network is the beneficiary.

HEAT GRADES

Air cooled exhaust	~35°C (95°F), heat nobody wants
Liquid capture	45 to 70°C (113 to 158°F), district heating grade
Digester hold	37°C (99°F) all year, a flat, forever load the pod serves directly

45 to 70°C

(113 TO 158°F) CAPTURED

~1.6 MW

TO YOUR NETWORK

Liquid pulls heat straight off the silicon.

It captures at temperatures your loop can use directly or with a small lift, instead of dumping low grade warmth to the sky.

About 1,500 sites can take a pod today

Roughly 700 district energy systems and 16,000 treatment plants in the country. You are early, not alone. Three tests, ten minutes with your own numbers.

~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	About 2 MW free continuously inside your contracted service, behind the meter
Heat sink	A hot water loop that takes warm water year round, or a small heat pump lift where it runs hotter
Fibre	A carrier or research drop on a concurrently maintainable path

Pass all three and the pod fits inside your fence line.

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 you are never the bottleneck.

Loop 01, off the silicon.

Cold plates pull heat straight off the chip, 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.

Loop 03, to you.

Heat goes to your loop, metered EN 1434 at your boundary. When your demand drops, the dry cooler rejects it instead.

DRY COOLER FIRST

The rejection plant is sized to carry the full load on its own, so the pod never needs your loop to stay cool. Take heat when you want it; hand it back when you do not. You are never on the hook for uptime.

One boundary, one meter.

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

Built, tested and cabled before it leaves the floor

A 2 MW pod arrives as a sealed, factory tested unit in two 13.716 m (45 ft) HC containers. 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
Heat returned	~1.6 MW to your 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. Craned onto its pad, connected at three points, computing within months.

Small inside

your fence line: two containers and a dry cooler bank, sized like a truck bay.

One pod. Four names on it.

The compute is contracted before the pod arrives, the unit is witness tested before it ships, and the operation is MicroLink's end to end.

THE NAMES

NVIDIA	GB300 NVL72 racks and the allocation
Vertiv	The powered and cooled pod, factory integrated and witness tested
Hydra Host	The compute demand and the offtake, contracted before arrival
MicroLink	The site, the host loop, the integration and the operation, end to end
IDEA	Channel partner, the International District Energy Association

You take no ownership and no operating role

YOU PROVIDE

Power	~2 MW at the agreed rate, behind the meter
Tie in	One exchanger into your heating loop or digester circuit
Path	A fibre route and a pad inside the fence

YOU RECEIVE

Heat	~1.6 MW recovered, 85% of delivered volume free, remainder at your avoided boiler cost
Fee	A site licence fee, plus 100 kW of compute if wanted

MICROLINK DOES EVERYTHING ELSE

Owns	Builds, funds, deploys and operates the pod
Carries	All technology, market and operating risk
Staffs	Remote NOC with regional hands on call
Meters	Heat at your boundary, EN 1434 class
Reports	ERE alongside PUE, every month

A good PUE says nothing about the heat

PUE measures how little a data center wastes running the building. ERE measures what happens to the energy afterwards. ERE is the number this company exists to move.

<0.5

ERE TARGET, YOUR NETWORK TAKING HEAT

1.12

DESIGN PUE, 1.09 TRI GENERATION

Reported together, always.

Both numbers come from EN 1434 class metering at your boundary. ERE below 1.0 means the site returns more usable energy than it loses as overhead.

YOUR MONTHLY HEAT STATEMENT

Delivered	MWh thermal, metered at your boundary
Displaced	Gas and emissions at your boiler efficiency
Reported	ERE alongside PUE, every month
Carries	A sustainability line for your annual report and your community consultation

The monthly statement is the sales instrument: yours to put in front of your board.

Prove one pod. Lock the design. Repeat it.

Rollout runs through your own sector bodies: IDEA, WEF, NACWA and the American Biogas Council. Specs are co published, so your engineers review a standard, not a vendor drawing.

Phase 01

Qualify

Gate A on the first site: headroom, heat sink, fibre and a clean permit path.

Weeks

Phase 02

First pod

Order, factory build, site prep in parallel, then connect and commission.

7 to 9 months, first of a kind

Phase 03

Framework

Design frozen with Vertiv and Hydra Host. Allocation and lead times set.

$12M per MW repeatable

Phase 04

Scale

Repeat pods across the qualified sites, each in under 20 weeks.

~250 site runway

Pilot cohorts deploy together, share commissioning learnings, and report to the same waitlist. Peers, not isolation.

The first 30 days

One agreement, one contact, one metered boundary. From signature to first compute is under 20 weeks, set by the factory slot, not by construction.

WEEK BY WEEK

Week 1	Site assessment against the three tests with your plant engineer: headroom, heat sink, fibre
Week 2	Interconnection spec issued against the co published association standard; tie in point agreed
Week 3	Power rate, heat terms and the site licence agreed; pad and route survey
Week 4	Factory slot confirmed with Vertiv; commissioning date set; single point of contact named

THE ASK

Host the pod. Take the heat.

Sign to host the pod and take the heat: provide ~2 MW at the agreed rate behind the meter, tie the pod into your heating loop at one exchanger, and receive ~1.6 MW of recovered heat, 85% of delivered volume free, plus 100 kW of compute if wanted. Join the waitlist at pods.microlinkdc.com; pilot cohort places are allocated in waitlist order within each channel.

~1.6 MW	85% free	$0	1 page
HEAT TO YOUR NETWORK	OF DELIVERED VOLUME	CAPITAL FROM YOU	FOR YOUR BOARD