


Operon Technologies

*Programmable power distribution
for GPU datacenters.*



The problem

GPU datacenters waste power converting and routing energy across multiple voltage rails. That lost power turns into heat, raises cost, and caps usable compute.

Today only 33 watts of every 100 from the grid reach the die as useful compute. The other 67 are dissipated in the cascade between the rack and the chip.

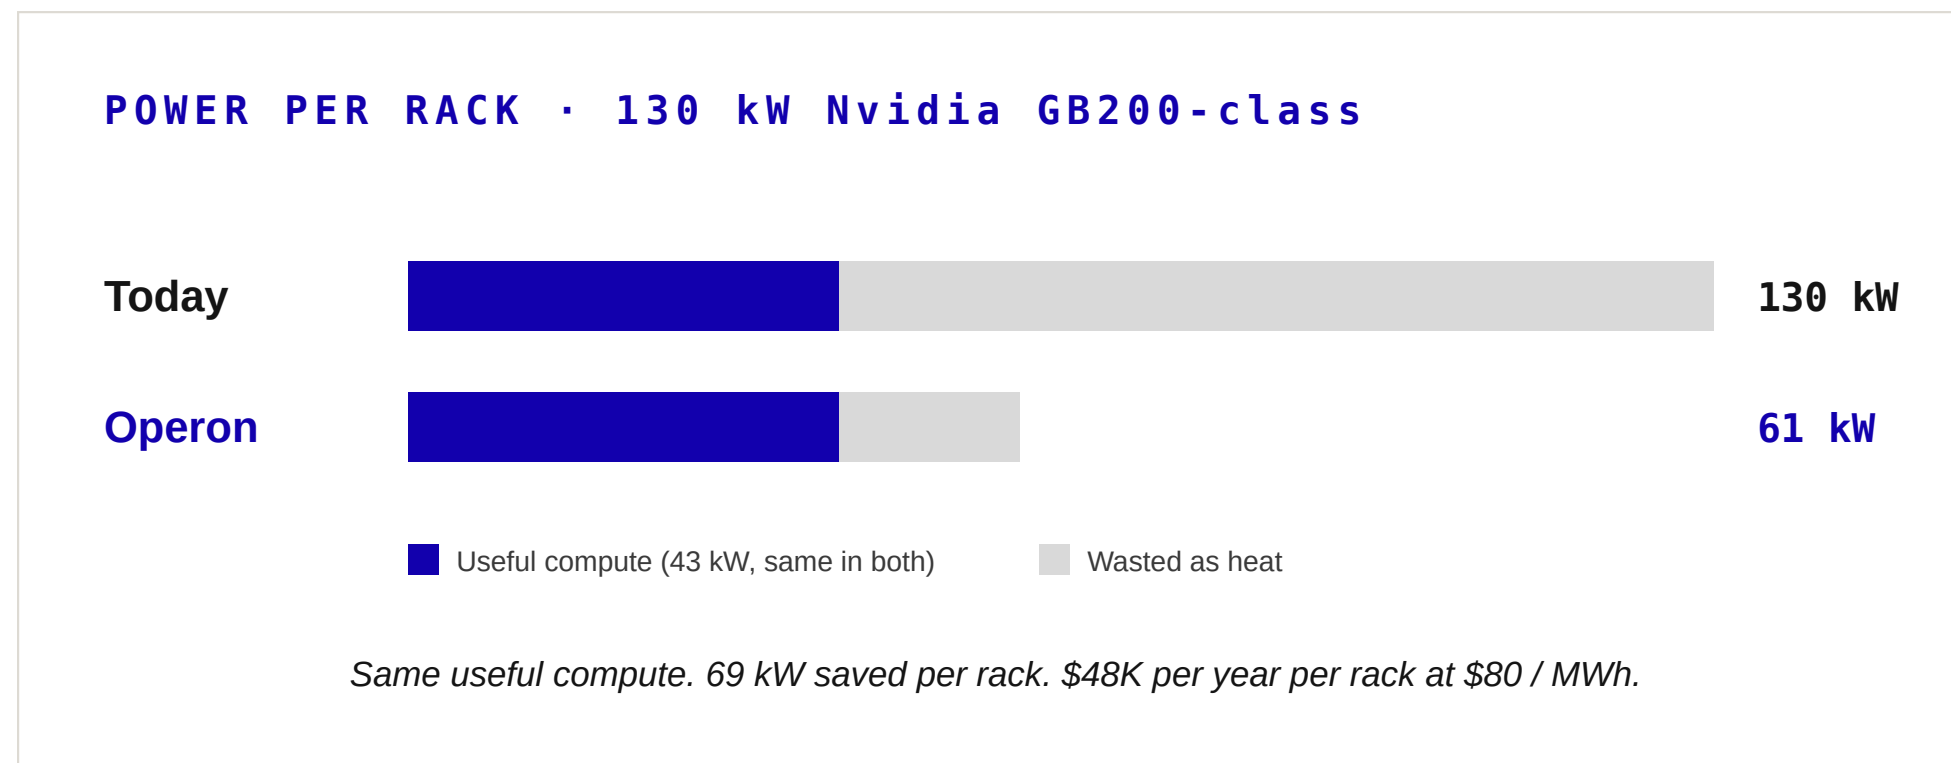
67%

*of AI grid power
lost before the die.*

*Vera Rubin draws 3,286 A at 0.7 V. Eleven times the
current load of Hopper three generations ago.*

The solution

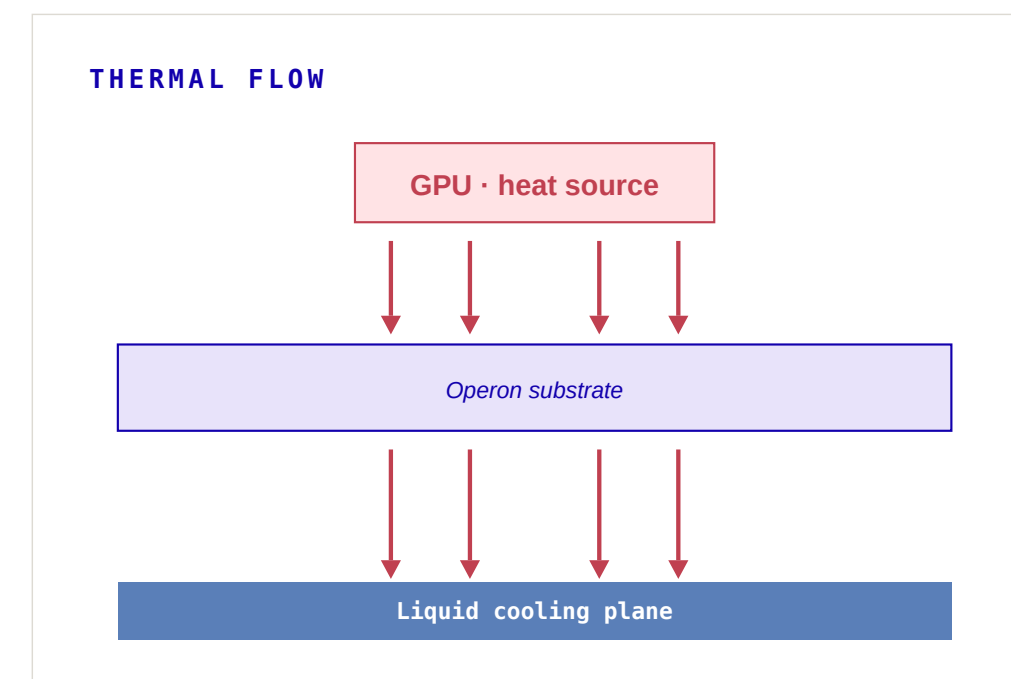
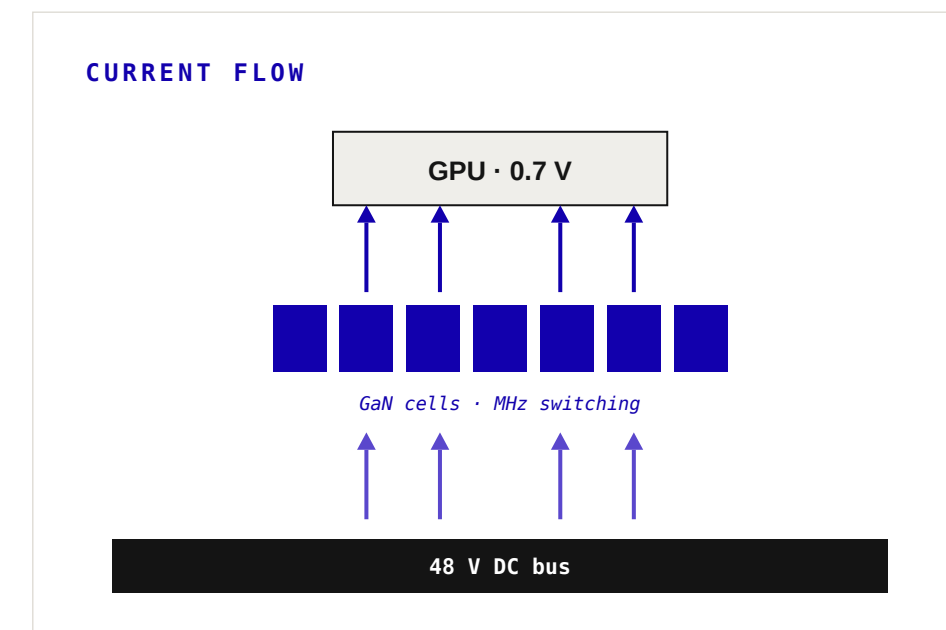
Operon is building a programmable power distribution layer that cuts conversion loss and heat while improving tokens per watt. Instead of hardwiring power for one chip generation, the system can adapt as chips and workloads change.



GaN switching cells, integrated magnetics, and vertical nano-interconnects, built directly into the motherboard PCB stack.

Operon power layer inside the motherboard PCB.

Signal layers, ground planes, and components stay conventional.



Why now

The bottleneck stopped being theoretical in the last twelve months.

FEB 2026

Intel discloses at ISSCC: at 5 kW per GPU, lateral power delivery converts only 42% of system input into compute. The problem is now quantified.

2025 - 26

Nvidia and the Open Compute Project move rack distribution to 800 V DC. Server architecture is being rewritten anyway.

NOW

Gallium nitride power transistors are commercially mature. The materials physics this architecture requires has finally arrived.

Competition

Nobody else moved to the substrate level

Player	Approach	Materials	Where it lives
Vicor / MPS / Infineon / TI	Discrete board-level VRMs	Silicon	On the motherboard
Empower → ADI (\$1.5B, May 2026)	System-in-package modules	Silicon	Next to the die
PowerLattice (\$25M Series A, Nov 2025)	Monolithic IVR chiplet	Silicon	Under the die
Operon	Materials-up substrate	GaN (wide-bandgap)	Inside the motherboard PCB

The incumbents stayed in silicon at the chip or board level. We moved up a layer with the materials physics that GaN enables.

Why We Will Win

Programmable power that survives the chip upgrade cycle.

Chip-agnostic substrate

Vera Rubin today, Vera Rubin Ultra in 2027, hyperscaler custom silicon after that. The substrate stays. Only the modules on top change.

Programmable current and voltage

Software-defined power per workload. The same hardware delivers different rails for inference, training, and idle. No board respin per chip generation.

No datacenter retrofit

Today every new GPU generation forces rack rework. With Operon, the rack stays. Datacenter operators plug and play.

Market

We sit under every chip in every AI server.

\$1T+

GLOBAL SEMICONDUCTOR SALES, 2026

AI is the engine making that number grow. Every compute die in an AI server sits on a motherboard. The motherboard needs a power layer. We are the power layer.

Buyers: hyperscaler custom-silicon teams and merchant server OEMs. Ten to fifteen globally.

What it's worth

Tokens per watt is the leverage point in AI compute economics. Operon doubles it.

\$7B

PER YEAR, PER GW

Additional token revenue at fixed grid power, at API price points used today.

\$350M

PER YEAR, PER GW

Electricity savings per gigawatt of original capacity, at fixed token throughput.

\$2.5B

PER GW

Avoided capex from cooling, power conversion, and supporting infrastructure no longer needed.

Technical risk

The hardest part of the build is also the moat.

THE OSAT BOTTLENECK

No commercial outsourced semiconductor assembly and test provider builds an integrated power substrate at panel scale today. Conventional OSATs (Amkor, ASE, JCET) operate at chip-package scale. PCB fabs (TTM, Compeq, Unimicron) make passive copper laminates. The Operon substrate sits between these two worlds.

MITIGATION

Hybrid manufacturing. A semiconductor partner fabricates the GaN switching layers and nano-via arrays at panel scale. A board house laminates the result into the multi-layer PCB stack. The coordination problem is the moat. Once one process line is qualified, replicating it across the industry is hard.

Team and ask

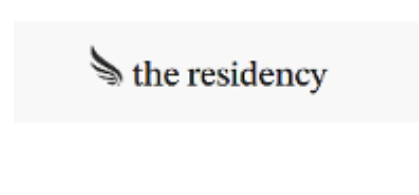
FOUNDER



Syrine Ben Driss

I've spent the last few years building across energy, biology, and applied infrastructure, from winning first place for a hydroelectric data center concept for Huawei, to co-designing a lab-on-chip with quantum dots for early cancer detection, to prototyping an IoT/AI system for pharma manufacturing, to building a fundraising app for a clinic in Sierra Leone.

IN PARTICIPATION



THE ASK

\$3M first round

18-month runway to the first substrate validation.

MILESTONES

- 2–3 Joint Development Agreements (JDAs)
- 1 paid pilot
- 1 manufacturing qualification partner
- 1 hyperscaler design-in evaluation
- 1 server OEM evaluation platform