



31 August – 1 September 2026

TALLINN, ESTONIA

Every Terabyte Is a Watt

Energy arithmetic for AI-scale storage

Alain Wilmouth CEO & Co-founder · 2CRSi · Æther Infrastructures



ABOUT 2CRSi

We build the machines this discussion is about

SINCE

2005

Founded in Strasbourg by two brothers and still headquartered there. Listed on Euronext Growth Paris.

COOLING

Three thermal paths

Air, direct-to-chip and immersion on one platform, so the customer picks the thermal strategy rather than the vendor.

ARCHITECTURE

OCP at 21 inches

OCtoPus: one shared 12 VDC power shelf instead of a PSU per node.
Labelled by Solar Impulse Foundation since 2021.

INSTALLED BASE

HPC and research

CERN, CEA, Caltech, Dassault Aviation, Airbus Defence, Naval Group and CGG among others.

We design and assemble the hardware.

That's why watts per terabyte is an engineering constraint here.

WHERE WE ARE

Three numbers that changed the storage business

330M TB

created every day, on a path to triple by 2030.

Six days now produce what all of 2010 produced.

Sold out

Western Digital's entire hard-drive output for calendar 2026.

Hyperscaler orders are locked into 2027 and 2028.

+55%

NAND contract prices, quarter on quarter, in Q1 2026.

Enterprise SSDs moved with them.

Capacity has stopped being a procurement problem.

It is now a physics problem, and physics is billed in kilowatt-hours.

THE SHIFT

The constraint moved three times in fifteen years

2010S

€ / TB

Buy the cheapest byte.

Floor space and racks are the limit.

2020S

W / TB

Buy the coolest byte.

The rack's power budget is the limit.

NOW

MW at the meter

Buy the byte you can power.


The grid connection is now the limit.

And storage watts are not billed once.

At a PUE of 1.2 to 1.56, every watt drawn by a drive pulls another 0.2 to 0.56 W of cooling behind it - before a single GPU is fed.

LEVER 1 — DENSITY

A hard drive draws 9 to 10 W whatever its capacity



W/TB at the drive. Sources: Western Digital (4 TB, 32 TB), Toshiba (10 TB, 24 TB). 44 TB figure derived at 9–10 W per drive.

So density is a free efficiency win

Moving from a 4 TB air drive to a 32 TB helium drive improves watts per terabyte roughly nine-fold — with no change to your software.

HAMR carries the curve further. Seagate's Mozaic 4+ ships 44 TB in volume to two hyperscalers today, 50 TB enters validation in 2027, and the published roadmap reaches a 100 TB drive.

Flash has done the same: 245 TB in a single E3.L slot, with 512 TB classes announced.

LEVERS 2 – 4

Density alone will not close the gap

2

Duty cycle

Host-managed SMR keeps writes sequential, so the drive stops rewriting itself. Power-disable and idle states let cold shelves stand down. Erasure coding avoids the RAID rebuilds that heat a whole array for days.

3

Data reduction

Compression and dedup multiply with density instead of competing with it. The cheapest byte is the one never written — it costs nothing to store, nothing to cool and nothing to replicate three times.

4

Placement

Tier by access temperature, not by department. Most AI training corpora are written once and read a handful of times; almost none of it belongs on flash after the first epoch.

Each lever is worth a fraction. Stacked, they are worth more than the next generation of media.

WHAT AI CHANGED

AI changed how the bytes are used

- **Training corpora**

Write once, read across epochs, then go cold. Enormous, and almost entirely misplaced on flash.

- **Checkpoints**

Bursty, synchronous, multi-terabyte. The cluster stops while they land — and an idle GPU is the most expensive idle in the building.

- **KV cache and inference**

Latency-critical, flash-resident, and the fastest-growing tier of all as context windows stretch.

- **Vector stores and RAG**

Small random reads at massive scale — an access pattern that barely existed in this industry five years ago.

Storage is now a GPU-utilisation problem

Accelerators are the scarcest and most power-hungry asset on the floor. Every point of utilisation lost to a stalled read is paid twice: once in capex, once in the watts burned idling.

Which is why storage is finally designed with the compute tier, not after it.

THE INFRASTRUCTURE LAYER

Efficiency you cannot buy from the drive vendor

Two-phase immersion

Up to 90%

less cooling energy than air, with PUE approaching 1.

2CRSi and Chemours signed a joint development agreement in February 2026.

OCP at 21 inches

Up to 23%

less power than equivalent 19-inch systems, through a shared power shelf instead of a PSU per node

DC Power Distribution

Fewer stages

between the transformer and the platter. Every conversion you remove is loss you never have to cool.

Storage-side BESS

40 MWh

between grid and datacenter, the batteries enable best cost energy storage to optimize opex.

THE OTHER HALF

87% of the footprint is use-phase — until you destroy it

90%

of decommissioned drives are physically destroyed, on the Circular Drive Initiative's estimate.

275×

the carbon difference between extending a device's life and retiring it early.

45.9 kg

CO₂e per year for a high-capacity HDD, 87% of it in the use phase, on Seagate's own LCA.

Shredding is now a choice, not a requirement.

IEEE 2883 and ISO/IEC 27040:2024 make cryptographic purge auditable and certifiable. In a market that is sold out through 2026, the drives already on your floor are the cheapest and cleanest capacity available anywhere.

ABOUT ÆTHER INFRASTRUCTURES

AI compute and decarbonised energy, in one company

What it is

A venture that develops the power first, then the site, then the compute — in that order, because that is the order in which the constraints actually bind.

Why it exists

Europe's ceiling on AI capacity is grid connection and lead time, not silicon. Power-first infrastructure has to be built, not retrofitted onto a shed.

Where

Two sites under development near Strasbourg, FR-SXB1 and FR-SXB2, on brownfield industrial land where the grid position, the water and the buildings already exist.


Who it serves

HPC, LLM training and industrial AI — with European operation treated as a design requirement rather than a label applied afterwards.

Storage physics runs all the way back to the substation. Somebody has to build that end of it.

AT GIGAFACTORY SCALE

Building the other end of the cable



Heat is an output, not a loss

Thermal valorisation into the local network turns the site's waste heat into a territorial asset — and turns a neighbour into a stakeholder.

A EuroHPC candidacy

Because European AI capacity that depends on somebody else's grid and somebody else's supply chain is not really capacity.

What it means for storage

When the site owns its power, cold tiers can be scheduled against the price of electricity rather than a flat tariff. Capacity becomes a dispatchable load.

WHAT TO MEASURE

PUE will not tell you whether your storage is efficient

- | | | |
|---|--|---|
| 1 | W/TB at rest and under load | Not the nameplate figure. Idle draw is where a capacity tier actually spends its life. |
| 2 | kWh per PB-year | Including facility overhead, so the cooling multiplier is visible instead of hidden in a site-wide average. |
| 3 | PB per rack per kW | The density that pays the bill, combining media, packaging and power architecture in one number. |
| 4 | kg CO₂e per PB | Embodied and use phase together. Reporting only one of them will flatter whichever choice you already made. |
| 5 | % of capacity reused at refresh | The only metric on this list that improves without buying anything. |

Ask every vendor in this building for those five. Most can answer the first.

TAKEAWAYS

Three things to take back

1

Density is the cheapest watt you will ever save

Nine times better W/TB from 4 TB to 32 TB, with no change to your software. HAMR and QLC extend that curve for another decade.

2

The rest of the gap closes above the drive

Duty cycle, data reduction, placement, cooling and power architecture. All of it ships against hardware you already own.

3

In a sold-out market, reuse is capacity

Ninety per cent of retired drives are destroyed. Sanitisation standards now make that a decision rather than an obligation.

Every terabyte is a watt. Design for the watt.



31 August – 1 September 2026

TALLINN, ESTONIA

Thank you

Questions and discussion

Alain Wilmouth CEO & Co-founder · 2CRSi

2crsi.com · Strasbourg, France

