HyperNext Data Centers | Research

HN-RP-015

The Rack Became the Computer

Preparing an Indian campus for rack-scale accelerated systems

What changes when the unit of purchase is a rack. Why the wide rack is not about power density. What the floor, the doors and the pipework have to be before the first one arrives.

SERIES	HyperNext Research
PAPER	HN-RP-015
ISSUED	1 February 2027
VERSION	1.0
CLASSIFICATION	Public release
CITATION	HyperNext Research, HN-RP-015

This is what's Next.®

The Rack Became the Computer

This paper is part of the HyperNext Research series. Method, assumptions and source data are stated so another operator can reproduce the analysis on their own facilities. Citation as "HyperNext Research, HN-RP-015" is welcome.

CONTENTS

01 What rack scale actually changes
02 The wide rack, and the arithmetic that surprised us
03 Feeding it at eight hundred volts
04 Weight, and the floor underneath
05 Taking the heat out
06 India first, and what that commits us to
07 References and standards

01 What rack scale actually changes

A data centre has always been a building that holds servers. The servers arrived in boxes, went into racks, and the rack was a piece of furniture: a frame, some rails, a pair of power strips. It was bought separately, often from a different supplier, and nobody thought of it as part of the computer.

Rack-scale systems end that. The accelerators, the interconnect between them, the power conversion and the liquid distribution arrive as one assembled machine, commissioned as one thing and warranted as one thing. The published AMD Helios platform puts seventy two accelerators into a single rack with thirty one terabytes of high bandwidth memory across them, which is a quantity of coupled silicon that only makes sense if it is built and tested together.

The consequence for the building is that the negotiable surface moves. Nobody is going to redesign a rack-scale machine to suit a hall. The hall accommodates the machine, at the dimensions, weight, supply and water conditions the machine specifies, and it does so with a lead time measured in construction seasons rather than in procurement cycles.

FIGURE 1

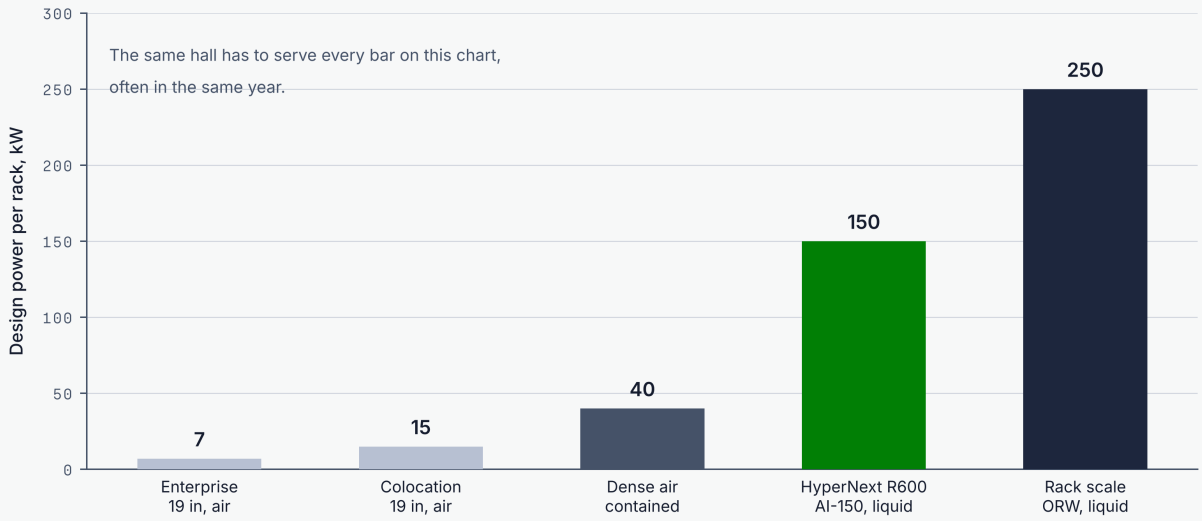


Figure 1. Design power per rack across the generations a hall has to serve. The gap between the first bar and the last is the design problem.

WHAT THIS PAPER IS NOT

HyperNext does not speak for AMD and does not publish measurements of other companies' equipment. Every figure here about a third party platform is a published figure, cited, and attributed to whoever published it. What is HyperNext's own is the design response: what we built, and why.

02

The wide rack, and the arithmetic that surprised us

The Open Rack Wide form factor is 1,200 mm across, twice the width of the rack it succeeds. The natural assumption is that this is how the industry gets more power into a square metre. We assumed it too, and the arithmetic does not support it.

In the space one wide rack occupies, two narrow racks fit. If those narrow racks are dense liquid-cooled units in their own right, the pair can draw more than the single wide rack does. On power per metre of row, the wide format loses. That is not a criticism of it, it is evidence that density was not what it was for.

What the width buys is a single enclosure large enough to hold one

FIGURE 2



Figure 2. The same length of row, populated two ways. Two narrow racks at 150 kW occupy the width of one wide rack at 250 kW and deliver more power into it.

scale-up domain without splitting it. Accelerators that have to talk to each other at memory speed do not want to be on opposite sides of a rack boundary, because every metre of interconnect and every connector between them costs latency, power and a place for something to go wrong. Keeping seventy two of them in one frame is an electrical and topological decision that happens to require a wider frame.

The second thing the width buys is room to work. A technician replacing a tray in a 600 mm rack packed to 150 kW is working in a space that was designed around the equipment with whatever was left over. At 1,200 mm there is room for service access to be designed rather than conceded, which for a machine that will be maintained for years is worth more than a few kilowatts per metre.

03 Feeding it at eight hundred volts

The form factor admits direct current from side power racks at low or high voltage. It does not tell an operator which to use, and the choice has consequences.

HyperNext feeds these racks at 800 volts direct current, taken straight from the campus distribution described in the earlier papers in this series, with no sidecar conversion rack standing beside them. This is worth stating carefully, because it is a HyperNext decision rather than something the published specification requires. The Open Compute

sidecar work specifies 400 volts as its nominal figure, with 800 named as an alternative. We standardised on the alternative.

The reason is the same one that runs through this whole architecture. At twice the voltage the current halves, and conduction loss falls with the square of current. For a machine drawing a quarter of a megawatt, the difference between feeding it at 400 volts and at 800 is a substantial quantity of copper and a measurable quantity of heat, both of which have to be paid for repeatedly over the life of the hall.

Removing the sidecar is the second half of it. A conversion rack beside every compute rack is floor area that computes nothing, a thermal load in the middle of the hall, and another assembly to maintain. Where the machine will take 800 volts directly, giving it 800 volts directly is the shorter path.

THE HONEST CAVEAT

Standardising above a published nominal means a thinner component market and insulation coordination that has to be done rather than inherited. HyperNext took that trade deliberately. An operator with a smaller estate, less in-house power electronics capability or a different component supply position might reasonably decide the other way, and would not be wrong to.

04

Weight, and the floor underneath

Everything above this point is electrical. The part that actually constrains the programme is structural, and it has to be finished first.

The published Open Rack Wide specification is rated to carry up to 3,500 kg of IT equipment. HyperNext plans on 4.5 to 5.5 tonnes for a delivered rack, and the gap between those two numbers is not a disagreement. The specification rates the payload. A building has to carry the payload plus the frame, the power shelves, the manifolds, the coolant in them and whatever else arrives bolted in. Designing the floor to the payload figure would be designing it to the wrong number.

That mass has to arrive somewhere and travel. On a multi-storey data centre it means goods lifts rated for it, a route from the loading bay that never passes over a slab designed for less, and door openings wide enough for a 1,200 mm rack with its packaging and its handling equipment. HyperNext specifies seven metre data hall door openings

for exactly this reason, which looks extravagant on a drawing and is simply the width that lets a machine in without dismantling it.

- The floor slab is designed for the loaded rack, not for the equipment rating, and not for an average spread across the hall.
- The delivery route is checked end to end as one path. A route is as strong as its weakest slab, and the weak one is usually a corridor nobody thought of as structure.
- Door and lift openings take the rack with its handling gear, because a machine that has to be partly dismantled to get in has to be partly reassembled by somebody who did not build it.
- Seismic restraint is designed at the loaded mass to IS 1893, which for a tall heavy rack is a different problem from restraining a conventional one.

None of this can be retrofitted. A slab is poured once, a lift shaft is sized once, and a door opening is formed when the wall is built. The equipment decision has a procurement lead time and the building decisions have a construction lead time, and the second is much the longer of the two.

05 Taking the heat out

A quarter of a megawatt in one enclosure cannot be cooled with air, and saying so is the easy part.

FIGURE 3

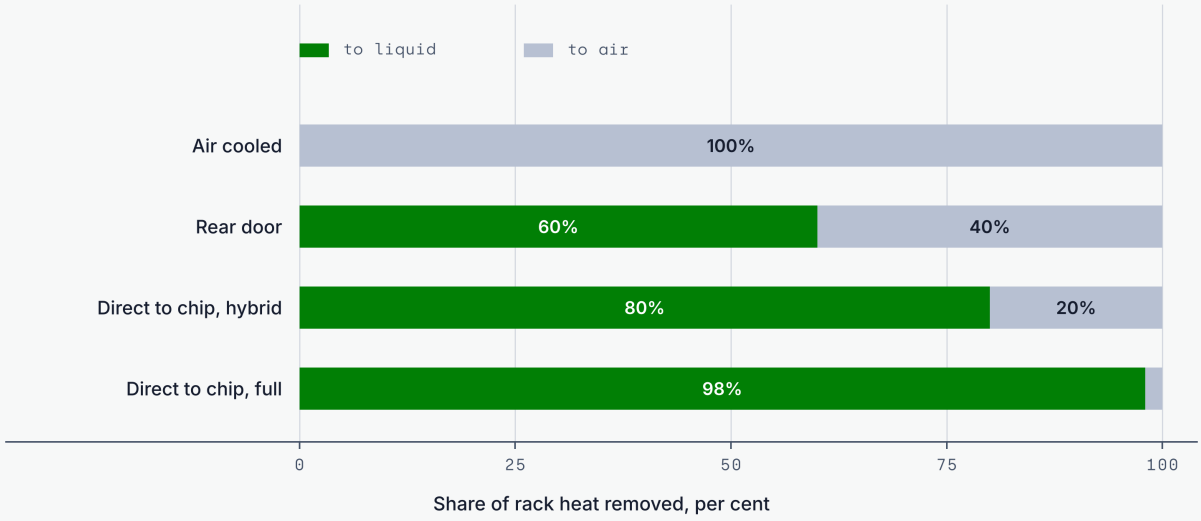


Figure 3. Where rack heat leaves, by cooling approach. The interesting column is the residual air fraction at the bottom, not the liquid fraction at the top.

HyperNext designs these halls for direct-to-chip liquid cooling taking essentially all of the rack heat. The HN-5 campus at Johor Bahru is built on that basis throughout. What this does not do is remove the air system, and the belief that it does is a recurring and expensive error.

A few per cent of the heat still leaves as air, from components that are not on the cold plate, from the power conversion, from losses in the distribution. A few per cent of a quarter of a megawatt is several kilowatts per rack, which in a row is a load that has to be collected and removed. Beyond that, the hall still needs ventilation for the people in it, it still needs a design condition for the equipment that is not liquid cooled, and it still has to behave sensibly when the liquid system is down for maintenance.

The right way to put it is that liquid cooling changes what the air system is for. It stops being the means of removing the computing load and becomes the means of conditioning the space. It gets much smaller and it does not get deleted, and a design that deleted it would discover the gap at commissioning.

06

India first, and what that commits us to

HyperNext brings rack-scale accelerated compute to India at HN-1 in Hyderabad in June 2027, at 12.5 MW, in a hall whose envelope can be raised to 14 MW if the demand is there.

Twelve and a half megawatts is a deliberate figure rather than a round one. It is what the hall can take with the cooling plant, the electrical distribution and the structure as designed, and it is stated as the envelope rather than as an ambition because the things that set it were poured and installed in advance. Raising it to fourteen is a change within the design, not a rebuild, which is the distinction that matters when a customer asks whether more is possible.

The larger commitment is HN-5 at Johor Bahru, 192 MW built for rack-scale systems end to end, fully liquid cooled, with a design basis of 80 per cent renewable supply, going live in July 2028. Ground was broken in November 2026. The two projects are different in scale and the same in kind: in both, the equipment date is the easy date and the civil work is the one that decides it.

TABLE 1

	HN-1 HYDERABAD	HN-5 JOHOR BAHRU
Rack-scale capacity	**12.5 MW, envelope to 14**	**192 MW**
Live	June 2027	July 2028
Cooling	Direct to chip	Direct to chip, campus wide
Supply	800 VDC, no sidecar	800 VDC, no sidecar
Renewable basis	Campus captive solar	80 per cent of supply

The two HyperNext campuses built for rack-scale systems. Dates are design dates at the time of issue.

The reason for saying India first out loud is not promotional. It is that a domestic campus able to take this class of machine, on an Indian floor slab built to Indian codes with Indian power electronics in the path, is the difference between Indian organisations renting this capability abroad and owning it at home. That was the reason HyperNext exists, and this is the first point at which it is demonstrable rather than asserted.

07 References and standards

01 **OCP Open Rack Wide (ORW) Base Specification v1.0.0**
Open Compute Project. The wide rack form factor contributed by Meta, including the 1,200 mm width and the 3,500 kg IT equipment rating cited in sections 2 and 4.

02 **AMD, Helios rack-scale platform, OCP Global Summit, October 2025**
Published platform description: 72 Instinct MI450 accelerators per rack, up to 432 GB of HBM4 per accelerator, 31 TB of HBM4 and 1.4 exaFLOPS of FP8 per rack. Figures are AMD's.

03 **OCP Diablo, v0.5.2, May 2025**
Open Compute Project sidecar power rack. Specifies 400 VDC nominal with 800 VDC as an alternative.

04 **ASHRAE TC 9.9, Thermal Guidelines for Liquid Cooled Data Processing Environments**
The facility water temperature classes referred to in section 5.

05 **IS 1893 (Part 1), Criteria for Earthquake Resistant Design of Structures**
Seismic basis for the structural work described in section 4.

06 **IS 875 (Part 2), Code of Practice for Design Loads, Imposed Loads**
The imposed floor loading basis for the data halls.

07 **HyperNext Research, HN-RP-012**
Breaking Direct Current at the Rack. Describes the protection these halls use.

HYPERNEXT RESEARCH

The Rack Became the Computer

Cite as: HyperNext Research, HN-RP-015, "The Rack Became the Computer", version 1.0, 1 February 2027. Available at hypernxt.com/research.

HyperNext Data Center Limited, CIN U63111TS2025PLC204792. Specifications of third party equipment are taken from published sources cited in the references and are the property of their owners. Nothing here is an offer of any product or service.