HyperNext Data Centers | Research

HN-RP-016

One Frame, Every Workload

The universal rack, and where the copper goes

Why four very different machines share one welded frame. What converting at the top of the rack saves. Why the heavy conductor should be short.

SERIES	HyperNext Research
PAPER	HN-RP-016
ISSUED	16 February 2027
VERSION	1.0
CLASSIFICATION	Public release
CITATION	HyperNext Research, HN-RP-016

This is what's Next.®

One Frame, Every Workload

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-016" is welcome.

CONTENTS

01 The problem with buying three rack families
02 One frame, four models
03 Where the copper goes
04 Converting at the top of the rack
05 Air today, liquid tomorrow
06 Built to Indian standards, and what that meant
07 The parts nobody specifies
08 References and standards

01

The problem with buying three rack families

Walk a HyperNext floor and you will pass a mainframe-class rack drawing under seven kilowatts, a colocation cabinet at twelve to eighteen, and a liquid-cooled accelerated compute rack at a hundred and fifty. They are not in separate buildings. They are often in the same row.

The conventional answer is to buy three rack families: a standard 19 inch cabinet for the servers, a vendor-specific enclosure for the mainframe equipment, and an Open Compute frame for the accelerated systems. Each comes from a different supplier, with different dimensions, different mounting, different doors and a different spares list.

That works until the floor changes. A hall built for storage wants to take compute. A tenant leaves and the next one brings different equipment. A refresh cycle doubles the power of the same workload in the same footprint. In each case the enclosures are wrong, and replacing a row of racks means emptying it, which in an operating facility is the expensive part.

HyperNext built one frame instead. Four models share a single welded structure, and what differs between them is the set of kits bolted to it. A rack converts from one model to another without cutting, welding or

FIGURE 1



Figure 1. The R600. One welded frame, 600 by 1,219 by 2,350 mm, with the kits for a given workload bolted to it.

re-coating, which means a hall can change what it does without changing what it is made of.

02 One frame, four models

The structure carries 2,200 kg on its levelling feet and rolls at 1,500 kg loaded, which it has to do because the rack ships with its equipment already in it.

FIGURE 2



Figure 2. The welded frame with its kit interfaces. Corner posts in IS 2062 E350, side rails on captive nuts, seismic anchor brackets at the base.

TABLE 1

MODEL	SERVES	MOUNTING	CAPACITY
AI-150	Accelerated compute, liquid	Open Rack V3, 44 OU at 48 mm	**150 kW**
IBM-50	Mainframe-class systems	EIA-310, 719 mm flange spacing	50 kVA per path
SV-25	General 19 inch servers	EIA-310, 43 RU	25 kVA per path
ST-25	Deep storage	EIA-310, deep frame at 1,600	25 kVA per path

Four models on one frame. Only the deep storage model changes the frame depth, and it does so because the equipment is physically longer rather than because the structure needed to differ.

Supporting two mounting standards in one frame is the part that costs something. Open Rack V3 works in 48 mm units across a 21 inch opening. EIA-310 works in 44.45 mm units across 19 inches. They are not compatible and no clever bracket makes them so, which is why the mounting is a kit: the frame carries the interfaces for both and one of them is fitted.

FIGURE 3

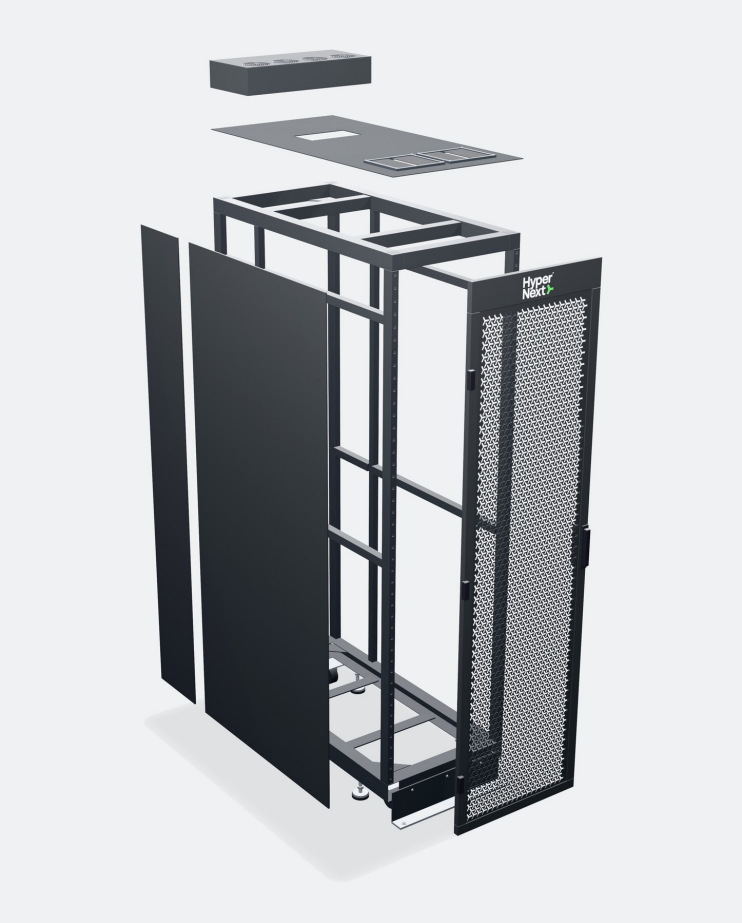


Figure 3. The rack as a kit of parts. Every model is the same frame with a different set fitted, at the factory or on site.

WHAT THIS RACK IS NOT FOR

Rack-scale accelerated systems in the Open Rack Wide format arrive as complete machines with their own frame, their own power shelves and their own distribution. The R600 is not placed in front of those and does not compete with them. It serves everything else on the floor, which on any real campus is most of the floor.

03 Where the copper goes

Copper is the quiet cost in a data centre. It is bought once, it is priced by a commodity market nobody controls, and the quantity needed is decided years earlier by a voltage choice that often gets made for other reasons.

The arithmetic is short enough to do here. Power is voltage times current, so at a fixed power the current falls in proportion to the voltage. Loss in a conductor is the current squared times its resistance, and resistance is inversely proportional to cross-sectional area. Hold the loss to a fixed fraction of the power delivered and the area required comes out proportional to the power, the resistivity and the length, divided by the loss fraction and the square of the voltage.

The voltage term is squared, and that is the whole story. Doubling the distribution voltage quarters the copper. Going from 50 volts to 800 is sixteen doublings worth of current reduction and a factor of 256 in conductor area.

FIGURE 4

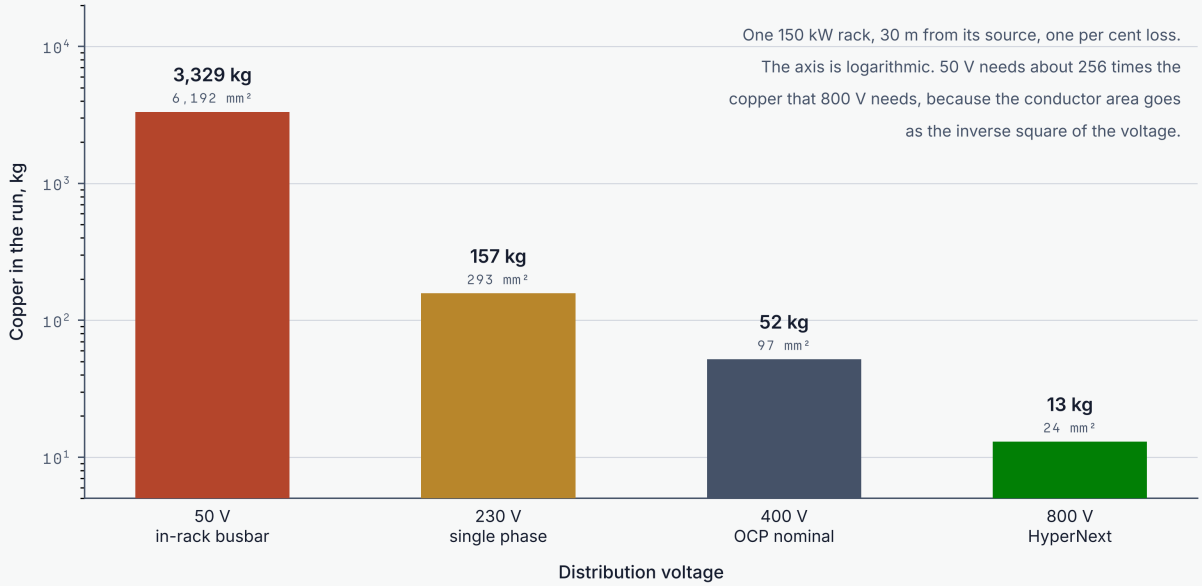


Figure 4. Copper in the run for one 150 kW rack thirty metres from its source, held to one per cent loss. The vertical axis is logarithmic.

Six thousand square millimetres of conductor to feed one rack is not a design, it is a reason not to do something. That figure is why low voltage direct current has never been distributed across a data centre floor, and it is worth stating plainly because the 800 volt argument is often made on efficiency when the stronger case is that the alternative does not physically fit in a cable tray.

04

Converting at the top of the rack

The conclusion people draw from that chart is to distribute at the highest voltage available. The more useful conclusion is about where the voltage changes.

Equipment does not run at 800 volts. The trays in an accelerated compute rack want about 50, and 150 kW at 50 volts is 3,200 amperes however the power got there. The high current is unavoidable. What is a design choice is how far it has to travel.

FIGURE 5

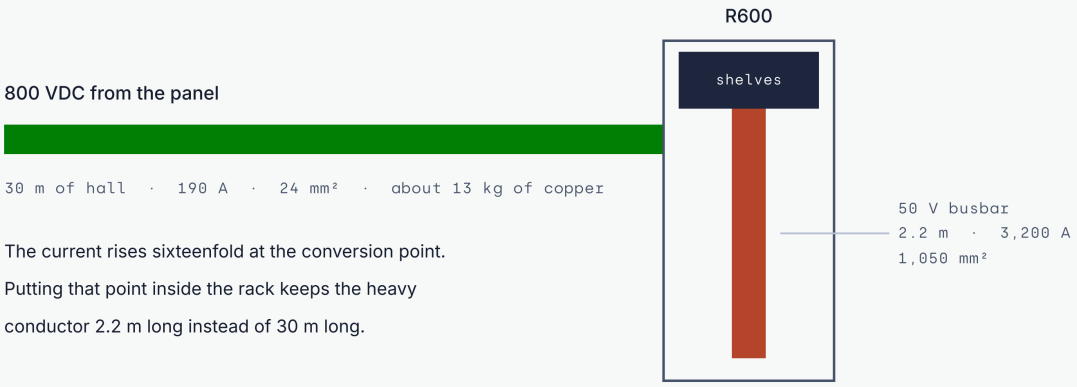


Figure 5. The same power, two conductors. The thin one crosses the hall, the thick one crosses the rack.

In the R600 the conversion shelves sit at the top of the rack rather than mid-rack or outside it. Eight hundred volts arrives at the roof on two independent paths, is converted in six shelves running five plus one, and leaves as a copper busbar down the back of the rack at 3,200 amperes. That busbar is six bars of 175 millimetres in electrolytic tough pitch copper, which is a substantial piece of metal, and it is 2.2 metres long.

Had the conversion happened at the end of the row, the same 3,200 amperes would have needed thirty metres of that cross-section instead of two. The busbar is heavy because it has to be. It is short because somebody decided where to put the shelves.

There is a second benefit that falls out of the same decision. Because the shelves are modular and the busbar is sized for the full 3,200 amperes from the start, adding capacity to a rack means adding a shelf. The rack does not have to be rebuilt, the busbar does not have to be replaced, and the hall does not have to be opened up. Capacity that has not been bought is not installed, and the room for it is already there.

FIGURE 6



Figure 6. Conversion shelves at the top of the rack. Two independent 800 VDC paths arrive at the roof and leave as the busbar below.

05

Air today, liquid tomorrow

Most racks delivered today will be air cooled. Most racks delivered in five years will not. The frame has to accept both without being replaced in between.

The model code carries a liquid level: L0 for air, L1 for direct liquid cooling. A rack moves from one to the other on site, without being moved. The full height kit fits two stainless manifolds with forty blind-mate ports, rated for 225 litres per minute at a ten kelvin rise. A smaller kit covers mixed racks where the lower part of the frame stays on air.

Every liquid joint sits over a drip tray, leak rope runs through every tray and pan to the rack controller, and the alarm threshold is five millilitres. That is a small volume deliberately: the point of detection is to find a weeping joint before it becomes a leak, not to confirm that a leak has

FIGURE 7



Figure 7. The liquid kit. Manifolds, valves, trays and leak detection bolt into the same frame that shipped air cooled.

happened.

Everything enters from above. That is partly a cleanliness decision and mostly a maintenance one: a service that can be isolated at the roof can be worked on without anybody lifting a floor tile or crawling behind a live rack.

06 Built to Indian standards, and what that meant

The rack is designed to Indian standards at the highest grade each offers, and manufactured accordingly. That was a constraint we chose and it had consequences.

- Structure in IS 2062 E350 C welded steel, with galvanised ledges, rather than the lighter sections a rack of this height could get away

FIGURE 8



Figure 8. The roof. Liquid, data and 800 VDC each take their own zone, so the floor stays clear and every service is isolated at the top.

with if it were only ever going to hold servers.

- Seismic restraint to IS 1893 Part 1, Zone V, top floor. Designing for the worst zone rather than the zone a particular campus sits in means one rack specification across the estate and no re-engineering when a campus is built somewhere else.
- Earthing to IS 3043, with a tinned copper frame earth bar and every touchable part bonded to 0.1 ohm maximum. Braided straps on every door and panel, because a door with a hinge for an earth path is a door with an intermittent earth path.
- Finish in zinc phosphate and fine texture powder coat at 80 micrometres minimum, with zinc flake fasteners rated to 480 hours of salt spray. Indian coastal campuses are the reason for the fastener specification, not the paint.

The doors carry the HyperNext mark as a perforation rather than as a badge: the brand prong, traced from the logo and punched 1,764 times, giving 56.5 per cent open area. The open area is the specification and the pattern is what satisfies it. A decorative perforation that restricted

airflow would be a bad door with a nice logo on it.

FIGURE 9



Figure 9. The door perforation. The pattern is the brand mark and the open area is the requirement it had to meet.

Racks ship loaded on a steel skid tied through the same brackets that later anchor the rack to the floor, at 1,500 kg, packed to a height that fits a forty foot high cube container. The skid, the anchors and the container dimension were designed together, which is the kind of thing that is invisible when it works.

07 The parts nobody specifies

A rack is bought on its capacity and lived with on its details. The items in this section appear on no comparison sheet and account for most of the difference between a cabinet that works and one that is merely correct.

Blanking is the clearest example. Any gap in the mounting plane lets hot

air from behind the equipment return to the front and be drawn in again, which raises inlet temperature without raising anything a monitoring system would notice until the equipment complains. The R600 ships with tool-less blanks in one and two unit heights and with air dams closing the side gaps that the 19 inch mounting standard leaves either side of the equipment. Both are cheap. Both are routinely left out, and a hall that left them out spends the difference on cooling for as long as it operates.

TABLE 2

DETAIL	WHAT IS FITTED	WHAT IT PREVENTS
Blanking	Tool-less 1 and 2 unit panels	Recirculation through the mounting plane
Side gaps	Air dams at the EIA gaps	Bypass around the equipment
Fans	**Four 120 mm units, three plus one**	**A single fan failure mattering**
Fan supply	2N from 800 VDC	The cooling sharing a fault with the load
Front cabling	150 mm frame, channels per unit	Cables crossing the air path
Monitoring	Rack controller reporting to DCIM	Finding out from the tenant

The details. Each is minor on its own and together they decide whether the thermal design that was drawn is the thermal design that runs.

The roof fan module is worth a sentence of its own because of how it is fed. Four units run in three plus one, so one can fail without the rack noticing, and they are supplied 2N from the 800 volt input rather than from a convenient outlet. A cooling system that shares a supply with the thing it is cooling fails at the same moment as the thing it is cooling, which is the wrong time.

Cabling gets a 150 mm front frame with finger channels at every mounting position, and a 300 mm side bay where a tenant needs more. This looks like a convenience and it is a thermal measure: cables routed across the front of equipment sit directly in its intake, and a rack that has been worked on for three years accumulates them until the air path is a negotiation.

The rack controller watches the fans and the leak chain and reports to the building monitoring system over a single connection. That is the whole specification and it is enough. The value is not in what it measures, it is that the measurement arrives somewhere a person is looking before the tenant telephones.

08

References and standards

01 **OCP Open Rack V3**
Open Compute Project. The 21 inch, 48 mm OU mounting standard followed by the AI model.

02 **EIA-310-E**
Cabinets, racks, panels and associated equipment. The 19 inch, 1.75 inch RU standard followed by the other three models.

03 **IS 2062, Hot Rolled Medium and High Tensile Structural Steel**
The E350 C grade used for the welded frame.

04 **IS 1893 (Part 1), Criteria for Earthquake Resistant Design of Structures**
Zone V, top floor, the seismic basis for the anchor brackets.

05 **IS 3043, Code of Practice for Earthing**
The basis for the frame earth bar and the bonding of every touchable part.

06 **IS 2551, Danger Notice Plates**
Hazard labelling on the enclosure.

07 **HyperNext Research, HN-RP-013**
Wide Bandgap Conversion at the Shelf. Describes the shelves at the top of this rack.

08 **HyperNext Research, HN-RP-015**
The Rack Became the Computer. Describes the rack-scale systems this rack does not serve.

HYPERNEXT RESEARCH

One Frame, Every Workload

Cite as: HyperNext Research, HN-RP-016, "One Frame, Every Workload", version 1.0, 16 February 2027. Available at hypernxt.com/research.

HyperNext Data Center Limited, CIN U63111TS2025PLC204792. This paper describes design decisions taken on HyperNext campuses. Figures are design values at the date of issue. Nothing here is an offer of any product or service.