HyperNext Data Centers | Research

HN-RP-013

Wide Bandgap Conversion at the Shelf

Gallium nitride at five hundred kilohertz, and what the frequency buys

Why switching faster makes a power shelf smaller. Where the gain is spent rather than banked. What a 45 degree inlet does to the whole argument.

SERIES	HyperNext Research
PAPER	HN-RP-013
ISSUED	2 January 2027
VERSION	1.0
CLASSIFICATION	Public release
CITATION	HyperNext Research, HN-RP-013

This is what's Next.®

Wide Bandgap Conversion at the Shelf

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-013" is welcome.

CONTENTS

01 What a power shelf has to do
02 Following Open Compute, and where we go past it
03 Silicon, silicon carbide, gallium nitride
04 The five hundred kilohertz decision
05 Where the frequency gain is spent
06 Designing for a forty-five degree inlet
07 Making it in India
08 References and standards

01 What a power shelf has to do

Eight hundred volts is a good voltage to move power at and a useless voltage to compute at. Something has to sit between the two, and in this architecture it is a shelf at the top of the rack rather than a room full of equipment at the end of the hall.

HyperNext builds two variants on one mechanical platform. The PS800-DD takes 800 volts direct current down to the roughly 50 volt bus that accelerated compute and modern server boards want. The PS800-DA takes the same input and produces 230 volt alternating current for equipment that has not moved, which on any real floor is a great deal of equipment and will remain so for years.

That second variant is the one that gets argued about internally, because it looks like a step backwards. The argument for it is simple arithmetic. A customer with a fleet of conventional servers is not going to replace them to suit our distribution choice. Either we convert for them somewhere, or we build a parallel alternating current infrastructure through the whole building to serve them. Converting in the rack, at the point of use, in a shelf that shares a platform with the one next to it, is the cheaper and the more flexible of those two.

FIGURE 1



Figure 1. The PS800 shelf. Open Compute form factor, six module positions, controller at the right, busbar clips at the rear.

TABLE 1

	PS800-DD	PS800-DA
Input	800 VDC	800 VDC
Output	Nominal 50 VDC	230 VAC
Module rating	16.5 kW	7.5 kVA
Shelf, N+1	**82.5 kW**	**37.5 kVA**
Efficiency	About 97.5% at full load	About 98.3% at the server
Serves	Accelerated and dense compute	Conventional servers and storage

The two variants. Shelf ratings are stated with one module held in reserve, which is how they are actually deployed rather than how they would be advertised.

02

Following Open Compute, and where we go past it

The shelves follow the Open Compute mechanical and management conventions in full. The dimensions, the busbar engagement, the management interface and the service access are not places to be original.

There is a real temptation in this industry to treat a published

specification as a starting point to improve on. It is almost always a mistake. The value of the specification is not that its dimensions are optimal, it is that everybody else uses them, so a shelf can be replaced, a spare can be stocked and a technician who has worked on somebody else's equipment already knows how this one comes out. Those are operational properties and they are worth more than a few millimetres of reclaimed depth.

Where HyperNext does depart is the input voltage. Open Compute's current sidecar work specifies 400 volts nominal, with 800 volts named as an alternative. The 400 volt choice is a sensible one and it was made largely to ride the electric vehicle component supply chain, which is a good reason that has nothing to do with data centre physics. HyperNext standardised at the higher figure because the current for a given power halves, and conduction loss falls with the square of current. The cost is a thinner component market and more careful insulation coordination, and we took that trade deliberately rather than by default.

FIGURE 2

Make-last, break-first on the power contacts.



Figure 2. Rear busbar engagement. The interlock sequence means the clip is never live while it is accessible.

One departure is not about the specification at all. At 800 volts a careless insertion is not a nuisance, it is an event. The shelves use a sequenced engagement so that the signal contacts make before and break after the power contacts, and the module cannot present a live connection at any point a hand can reach. This is ordinary practice at this voltage in other industries and it is new to a lot of data centre hardware.

03

Silicon, silicon carbide, gallium nitride

Three materials are available for this job and the choice between them is not a matter of one being better. They occupy different places and the places overlap.

Silicon is mature, cheap, available from everybody and has been engineered against its own limits for fifty years. Those limits are set by its critical electric field, which is about 0.3 megavolts per centimetre. A device that has to block a given voltage needs a drift region thick enough not to break down, and the thickness of that region sets the resistance the device has when it is on. For silicon at 800 volts that resistance is high enough to hurt.

Silicon carbide has a critical field near ten times higher, so the drift region can be roughly ten times thinner for the same blocking voltage. The relationship between blocking voltage and the best achievable specific on-resistance scales with the cube of the critical field, which is why the improvement is so much larger than the ratio of the fields suggests. Silicon carbide is the right material for the kilovolt and multi-kilovolt work described in the previous paper in this series.

Gallium nitride has a comparable critical field and a quite different structure. The devices used here are lateral rather than vertical, which caps the voltage they reach comfortably but gives them very small gate charge and, in the configuration we use, essentially no reverse recovery charge at all.

REVERSE RECOVERY IS THE THING THAT MATTERED

When a silicon body diode conducts and is then forced off, stored charge has to be swept out before it blocks, and every one of those events costs energy. Since the cost is paid per switching cycle, it is the term that punishes high frequency hardest. Remove it and the economics of frequency change completely. That, more than any headline efficiency figure, is why the shelf uses gallium nitride.

FIGURE 3

The minimum is broad, which is why the choice is a judgement rather than a calculation.

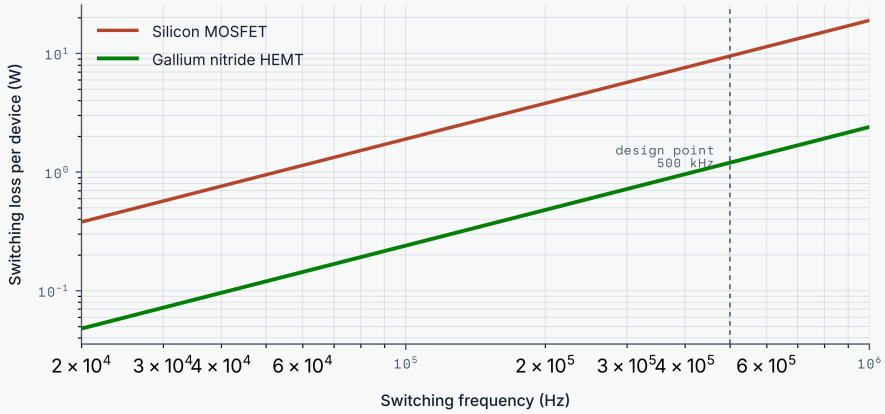


Figure 3. Loss components against switching frequency for the same converter duty. Conduction loss is flat with frequency, switching loss rises with it, and magnetic loss falls as the components shrink.

04

The five hundred kilohertz decision

The reason to switch faster is magnetics. It is almost the only reason, and it is enough.

A transformer transfers a fixed amount of energy per cycle for a given core and winding. Run it at twice the frequency and it transfers twice the power through the same iron, or the same power through half of it. The relationship is not exactly linear once core loss and winding effects are accounted for, but the direction is not in doubt, and magnetic components dominate the volume and the mass of any isolated converter. Shrink them and the shelf shrinks.

FIGURE 4

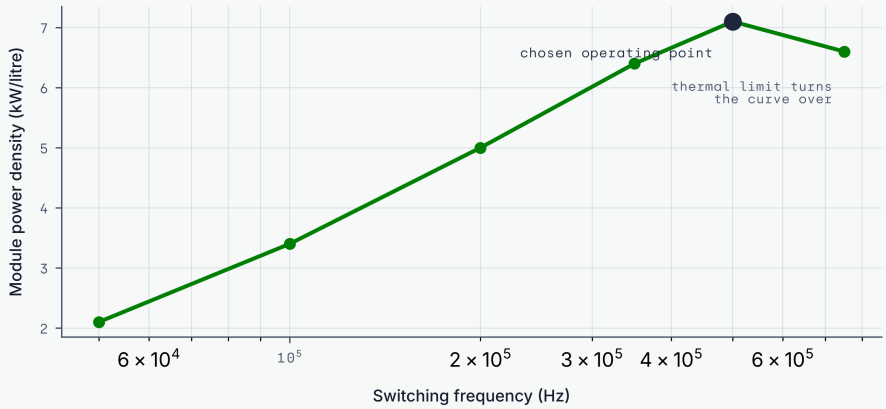


Figure 4. Indicative magnetic component volume against switching frequency at constant throughput, normalised to a 50 kilohertz baseline.

So why not go faster still? Because every term that scales with frequency is working against you at the same time. Switching loss rises, gate drive power rises, the parasitic inductance of the layout stops

being a detail and becomes the dominant constraint on the switching node, and electromagnetic emissions move up into bands that are harder to filter and harder to certify. Core loss itself rises steeply with frequency, which partly offsets the volume saving that motivated the whole exercise.

Near five hundred kilohertz the magnetics have shrunk by most of what is available, gate drive is still manageable, layout is difficult but tractable on a board that can be manufactured at volume, and the emissions problem is hard rather than unsolved. The curve in Figure 2 is flat enough around its minimum that a nearby frequency would also have worked. What would not have worked is treating the choice as a number to maximise.

500 kHz AND NOT HIGHER

The frequency was chosen where the magnetic saving had mostly been collected and the costs it creates were still ones we could engineer against rather than design around.

05

Where the frequency gain is spent

An efficiency figure is only interesting if it says where it was measured and at what load. Most published figures do neither.

FIGURE 5

Measured at the shelf output, not at the module.

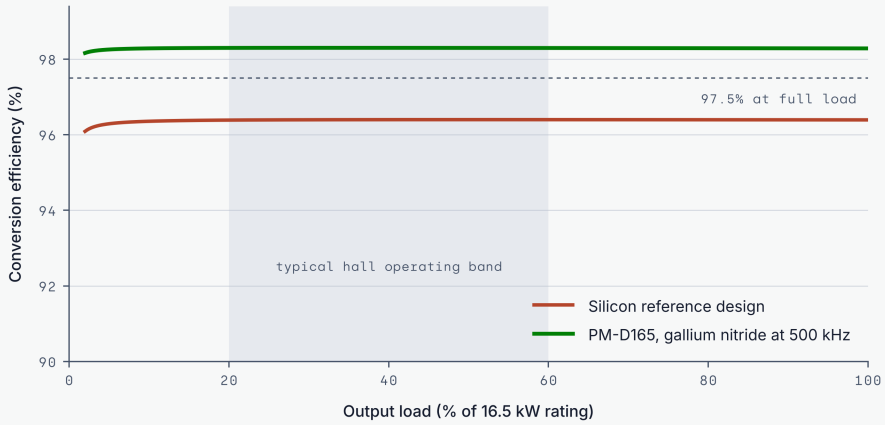


Figure 5. Efficiency against load for the DC-DC shelf. The curve peaks below full load, which is where equipment actually spends its life.

The DC-DC shelf reaches about 97.5 per cent at full load, with a peak above that at partial load. The shape matters more than the peak. A

shelf sized with one module in reserve and a workload that is not pinned at maximum spends most of its hours somewhere in the middle of that curve, which is the part of it worth designing for. A converter tuned for a headline figure at full load and allowed to fall away below it will lose that advantage every hour it runs.

The DC-AC shelf reaches about 98.3 per cent measured at the server. That boundary is the honest one to quote, because it is where the power is consumed and because it includes everything between the shelf and the load. A figure quoted at the shelf terminals excludes losses that someone still pays for.

It is worth being clear about what a point of efficiency is worth, because the number sounds small. At campus scale the saved power is not the whole of it. Loss that does not happen is loss that does not have to be removed as heat, so the cooling plant sized to remove it does not have to exist, and the capital it would have consumed is available for something that computes. That second effect is the larger of the two and it is the one that is usually left out of the comparison.

FIGURE 6

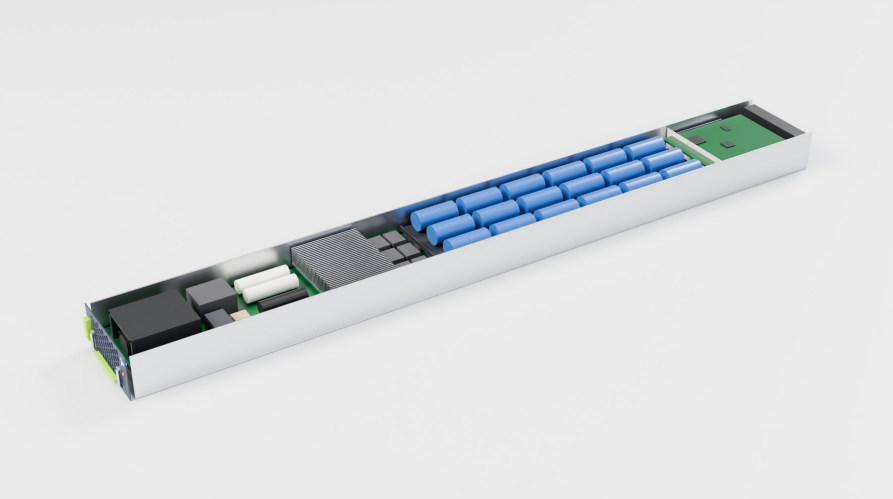


Figure 6. A conversion module. The magnetics are the components the frequency decision was made to shrink.

06 Designing for a forty-five degree inlet

The shelves sit in white space and take cold aisle air. They are designed for an inlet below 45 degrees, and that single figure constrains more of the design than anything else in this paper.

A converter designed for a 25 degree inlet has a comfortable margin between ambient and the junction temperature its devices can tolerate, and that margin can be spent on cheaper parts, smaller heatsinks or a quieter fan curve. Raising the inlet to 45 degrees removes twenty degrees of that margin before the design starts. Everything downstream pays: the device has to be selected for a higher junction temperature at rated power, the thermal path from die to air has to be shorter, and the board area given to heat spreading cannot be reclaimed for components.

The reason to accept that is what it does to the rest of the facility. A hall that can run warm needs less mechanical cooling, can use free cooling for more hours of the year, and in the Indian climates where HyperNext builds that difference is a large fraction of the annual cooling energy. The power electronics give up margin so that the building does not have to.

FIGURE 7

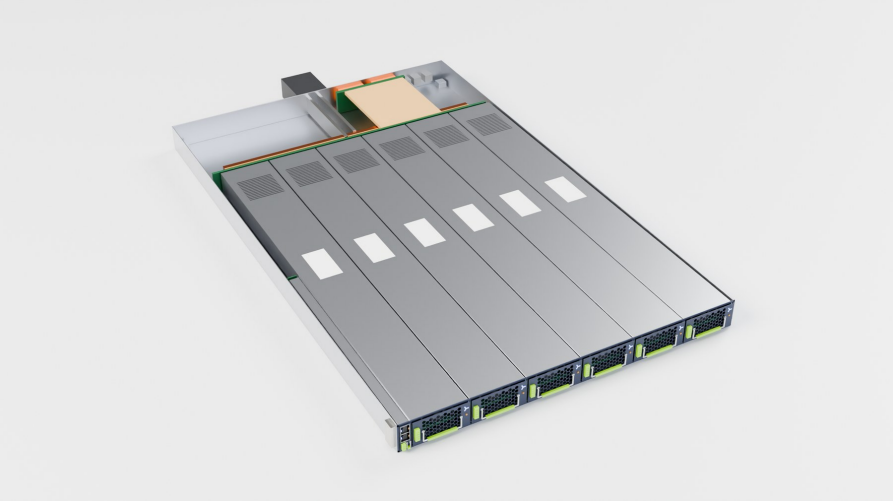


Figure 7. Shelf interior with the lid removed. The magnetics and the semiconductors sit on separate thermal paths rather than one behind the other.

Putting the hot components in series in one air path is the error that is easy to make and hard to correct later. Air that has already passed over the magnetics arrives at the semiconductors warmer than the inlet, which eats the margin that was just bought at some expense. Separating the paths costs mechanical complexity and it is the right trade at this inlet temperature.

FIGURE 8



Figure 8. Module removal under load. The shelf holds its output through the exchange on the remaining modules.

07

Making it in India

Where the shelves are made, and how the work is divided.

Bare board fabrication and assembly are outsourced to specialist suppliers, because a high frequency power board has layout and process requirements that a general purpose line does not meet, and because that industry exists and is good at this. Bringing it in-house would mean learning a mature discipline badly for the sake of owning it.

Firmware is developed in parallel with the hardware rather than after it. That is the only arrangement that catches a control problem while the board can still be changed, and it is the part of a power electronics programme most often sequenced the wrong way round. A converter whose control loop is written against finished hardware inherits every decision that hardware already made.

Final assembly, test and burn-in are done in-house in India. Keeping those three together matters more than where any one of them happens: the people who assemble a shelf see what comes back from test, and a fault found at burn-in reaches the bench that built it rather than a report that crosses a border.

Electromagnetic compatibility is the discipline this frequency makes expensive. At five hundred kilohertz the emission spectrum sits where filtering is harder, and a rack of identical units behaves differently from one unit on a bench. Compliance is demonstrated per unit against IEC 61000-6-4, and the aggregate behaviour is something the design anticipates rather than discovers.

08

References and standards

01 **OCP Diablo, v0.5.2, May 2025**
Open Compute Project sidecar power rack specification. Specifies plus or minus 400 VDC nominal with 800 VDC as an alternative. HyperNext standardised above the nominal.

02 **OCP ORv3 power shelf specification**
The mechanical and management conventions the PS800 shelves follow in full.

03 **IEC 62477-1**
Safety requirements for power electronic converter systems and equipment.

04 **IEC 61000-6-4**
Electromagnetic compatibility, emission standard for industrial environments. The constraint that grows with switching frequency.

05 **B. J. Baliga, Fundamentals of Power Semiconductor Devices, 2nd ed., Springer, 2019**
The figure of merit relating blocking voltage to specific on-resistance that section 3 rests on.

06 **ASHRAE TC 9.9, Thermal Guidelines for Data Processing Environments, 5th ed.**
The inlet temperature classes referred to in section 6.

07 **HyperNext Research, HN-RP-012**
Breaking Direct Current at the Rack. Describes the panel that feeds these shelves.

HYPERNEXT RESEARCH

Wide Bandgap Conversion at the Shelf

Cite as: HyperNext Research, HN-RP-013, "Wide Bandgap Conversion at the Shelf", version 1.0, 2 January 2027. Available at hypernxt.com/research.

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