HyperNext Data Centers | Research

HN-RP-011

Medium-Voltage Solid-State Transformation

33 kV to the rack without a distribution transformer

Topology and device physics. Why the magnetics shrink. How a direct current network clears a fault. What the efficiency claim actually counts.

SERIES	HyperNext Research
PAPER	HN-RP-011
ISSUED	1 December 2026
VERSION	1.0
CLASSIFICATION	Public release
CITATION	HyperNext Research, HN-RP-011

This is what's Next.®

Medium-Voltage Solid-State Transformation

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-011" is welcome.

CONTENTS

01 The chain an 800 VDC campus does not build
02 What a solid-state transformer is, and what it is not
03 Device physics: why silicon carbide and not silicon
04 Clearing a fault without a current zero
05 The efficiency account, stage by stage
06 Copper, iron and floor area
07 References and standards

01

The chain an 800 VDC campus does not build

A conventional data centre converts its incoming supply up to six times before a processor can use it. Each conversion is a separate asset that has to be bought, housed, cooled, maintained and eventually replaced, and each one takes a cut of the power on the way past.

SEE FIG 1

Figure 1 sets the two chains against each other at the same scale.

The chain runs from the incoming medium-voltage supply through a distribution transformer to low-voltage switchgear, then through the uninterruptible supply, which rectifies to a direct-current link, holds a battery across it and inverts back to alternating current. A power distribution unit transforms again. The rack power supply rectifies one last time to feed the board. Six conversions, four of which exist only to return the power to the form it started in.

HyperNext campuses are built on a different chain. Thirty-three kilovolts reaches a medium-voltage solid-state transformer, which delivers 800 volts direct current to a busway. The busway runs to the rack, where a power shelf steps down to the board. There is no distribution transformer, no low-voltage switchgear lineup, no central uninterruptible supply and no battery plant. Ride-through sits at the rack instead.

FIGURE 1

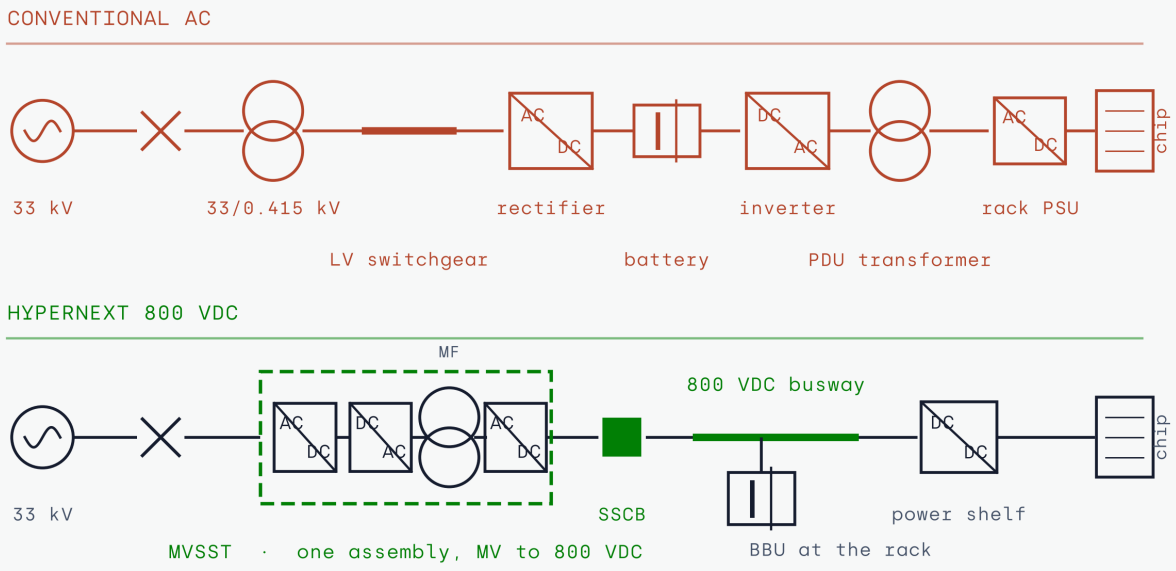


Figure 1. The two power chains, drawn to IEC 60617 symbols. Conventional above, HyperNext below. Both start at 33 kV and end at the same processor. Open cross is a mechanical breaker, filled square is solid state.

~95% REACHES THE CHIP

End to end across the HyperNext chain, at the design point. The conventional chain delivers roughly 75 per cent of what the meter records.

The difference is paid twice. Power lost in a conversion stage leaves as heat, and that heat is then removed by a cooling plant which itself draws power. A percentage point of conversion loss therefore costs more than a percentage point of the electricity bill, and the second payment is what couples the power chain to the cooling design. Section 5 puts numbers against both.

The absences matter as much as the efficiency. A central battery plant is the largest single fire load in a conventional facility, and it carries suppression, containment, ventilation and insurance with it. None of that is designed in here, so none of it has to be removed later.

02

What a solid-state transformer is, and what it is not

A solid-state transformer is not a transformer with electronics bolted to it. It is a power-electronic converter that happens to provide the two things a transformer provides: a voltage ratio and galvanic isolation.

The medium-voltage alternating supply is rectified to a direct-current link. That link drives an inverter running at several kilohertz, which feeds an isolation transformer. The transformer output is rectified again to the 800 volt bus. Three conversions happen inside the assembly, and it is worth saying so plainly, because the count is where most published comparisons become slippery.

ON COUNTING CONVERSIONS

HyperNext states two conversion stages against a conventional six. That count is at system level: one transformation from medium voltage to the 800 volt bus, and one at the rack. Inside the solid-state transformer there are three internal stages, as Figure 1 shows. The honest comparison is not the stage count but the measured end-to-end efficiency, because stages at medium frequency inside one assembly are far cheaper in loss than the discrete assets they replace.

The gain comes from frequency. A transformer core is sized by the volt-seconds it must support between flux reversals, so for a given throughput the required core area falls roughly in inverse proportion to the switching frequency. Moving from 50 hertz to a few kilohertz takes an order of magnitude out of the iron and the copper wound around it. That is the whole reason the assembly fits where a distribution transformer and its bund would not.

What the assembly must still do is everything the distribution transformer did, and several things it could not. It holds galvanic isolation through the medium-frequency core. It regulates output voltage, which a passive transformer never did. It limits fault current by control rather than by winding impedance. It runs power in both directions under control, which matters as campuses add storage and on-site generation.

FIGURE 2

Core area product scales as 1/f. Mass follows volume, so it scales more slowly than area.

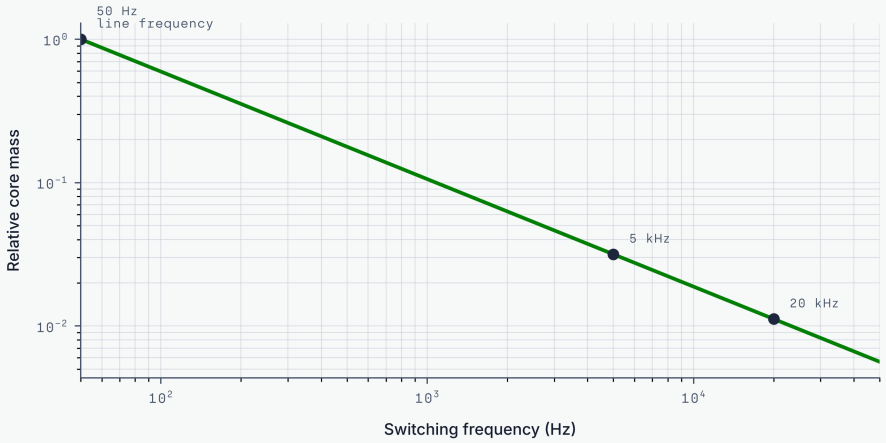


Figure 2. Relative core mass against switching frequency at constant throughput. This is a textbook relationship rather than a HyperNext measurement: the required core area product scales as the inverse of frequency, and mass follows volume.

TABLE 1

FUNCTION	DISTRIBUTION TRANSFORMER	MVSST
Voltage ratio	Fixed by turns ratio	Regulated under control
Galvanic isolation	At line frequency	At medium frequency
Fault current limiting	By winding impedance	By control
Bidirectional	Passive	Controlled
Output form	Alternating current	**Direct current, to the busway**
Relative core mass	1.0	**About 0.1**
Failure mode	Slow, thermal	Fast, electronic

What carries over from the asset being replaced, and what does not. Core mass follows the relationship plotted in Figure 2. The change in failure mode is not a footnote: electronics fail faster and more completely than iron, which is why redundancy is designed at the assembly level.

03

Device physics: why silicon carbide and not silicon

Silicon carbide is not an incremental improvement on silicon for this duty. It is what moves the architecture from a laboratory demonstration to something that can be specified, bought and maintained.

The property that matters is the critical electric field, the field a material will hold before it breaks down. For 4H silicon carbide it is near 2.8 megavolts per centimetre against roughly 0.3 for silicon, a factor close to ten. A device blocks voltage across a lightly doped drift region, and the thickness of that region is set by how much field the material

tolerates. Ten times the field means about a tenth of the thickness at the same blocking voltage, with roughly ten times the doping.

Both changes push the same way. The resistance of the drift region falls with its thickness and with its doping, so the specific on-resistance of a silicon carbide device sits two to three orders of magnitude below the silicon limit at the voltages a medium-voltage converter cell needs. Figure 3 plots the unipolar limit for both materials across the range.

FIGURE 3

Ron,sp rises with the square of blocking voltage and falls with the cube of critical field. The cube is where the advantage sits.

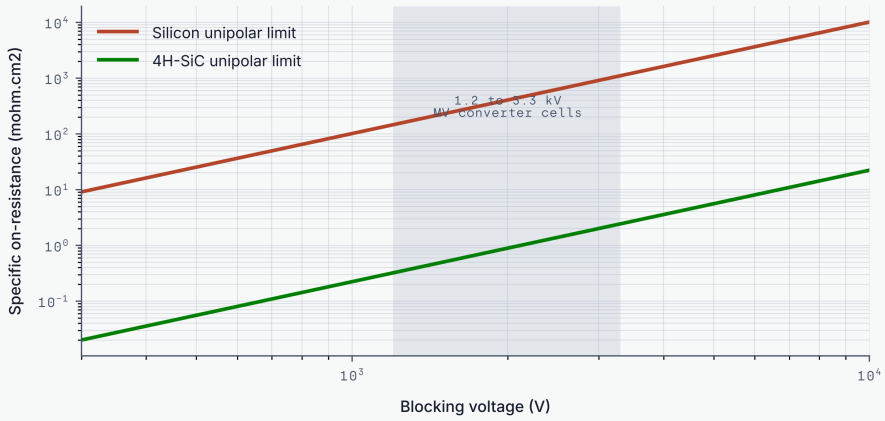


Figure 3. Unipolar specific on-resistance limits for silicon and 4H silicon carbide against blocking voltage, computed from the standard relationship using published material constants. The shaded band is the 1.2 to 3.3 kV class used in medium-voltage converter cells.

Two further properties decide whether the device survives in service. Silicon carbide conducts heat about three times better than silicon, so the die sheds heat into its package more readily at the same flux. And because these are majority carrier devices, there is no stored charge to sweep out at turn-off and no tail current, which is what allows switching in the kilohertz range at medium voltage without the losses becoming the dominant term.

The honest side of the ledger. Silicon carbide costs more per ampere than silicon and will for some time. Gate oxide reliability under high field and high temperature remains the subject of active work and is the reason gate drive margins are set conservatively. The fast edges that make the device efficient also stress winding insulation and raise electromagnetic compatibility work from an afterthought to a design discipline. The supply base for medium-voltage modules is thin, and that is a procurement risk rather than an engineering one.

04

Clearing a fault without a current zero

An alternating current circuit offers a free opportunity to interrupt itself one hundred times a second. Direct current offers none. This single fact is the largest engineering difference between the two architectures, and it is not widely appreciated outside the people who design the protection.

A mechanical alternating current breaker separates its contacts, draws an arc, and waits for the current to pass through zero on its own. The arc extinguishes at that point and the gap recovers. Clearing takes tens of milliseconds and the device can be comparatively simple because the waveform does most of the work.

In a direct current network the current does not pass through zero. It has to be forced there, or the energy in the arc has to be dissipated until the arc can no longer sustain itself. Mechanical direct current breakers do this with a commutation circuit that injects a counter-current, and they are large, slow and expensive by comparison. A solid-state breaker takes the other route: a semiconductor switch turns the current off in microseconds, with no arc at all.

FIGURE 4

Damage at a fault scales with I squared t. Three orders of magnitude is the difference between a replaced module and a replaced lineup.



Figure 4. Let-through energy for three interrupting technologies at a common 20 kA prospective fault, computed as the square of current times clearing time. The scale is logarithmic. Clearing time is the whole argument.

The consequence is not only speed. Damage done at a fault scales with the square of the current and with the time it is allowed to flow, so a device clearing in ten microseconds instead of twenty milliseconds lets through roughly three orders of magnitude less energy. What would have been a destroyed assembly becomes a replaced module. Arc flash incident energy scales with arcing time in the same way, and a fault that is never allowed to establish an arc changes the protective equipment

the maintenance team has to wear.

What is harder is selectivity. In an alternating current network, coordination between an upstream and a downstream device is arranged on time-current curves, using the milliseconds between them. Microsecond devices have no such room. Coordination is instead arranged on current thresholds and on the rate of rise of current, and the margins are narrow enough that they have to be proven by simulation before they are proven on site. Section 7 says what we had to build to do that.

THE ENERGY HAS TO GO SOMEWHERE

Turning off a current in microseconds does not remove the energy stored in the inductance of the cable run. That energy appears as a voltage across the opening device and has to be clamped and absorbed, which is what the snubber and the surge arrester across a solid-state breaker are for. A specification that quotes interruption time without stating the clamping energy is incomplete.

05

The efficiency account, stage by stage

Published efficiency comparisons in this industry are often a single number with no account behind it. The figures below are design-point values at the stated loading, stage by stage, so the arithmetic can be checked and argued with.

TABLE 2

STAGE, CONVENTIONAL AC	STAGE EFFICIENCY	CUMULATIVE
Incoming at 33 kV	100.0%	100.0%
Distribution transformer	99.0%	99.0%
LV switchgear and distribution	99.8%	98.8%
UPS rectifier	97.5%	96.3%
UPS inverter	97.0%	93.4%
PDU transformer	98.5%	92.0%
Rack power supply	92.0%	84.7%

Conventional chain at the design point. Cumulative efficiency is the running product. The figure is above the roughly 75 per cent quoted for a working facility because this table excludes distribution losses in cable runs and assumes every stage sits at its best operating point, which no real hall does.

TABLE 3

STAGE, HYPERNEXT 800 VDC	STAGE EFFICIENCY	CUMULATIVE
Incoming at 33 kV	100.0%	100.0%
MVSST, 33 kV to 800 VDC	97.5%	97.5%
800 VDC busway	99.7%	97.2%
Rack power shelf	98.0%	95.3%

HyperNext chain on the same basis. Four rows against seven. The shelf figure is the DC to DC variant; the DC to AC variant used for alternating current servers runs close to 98.3 per cent at the server.

FIGURE 5

The gap at the right edge is what the architecture is for.

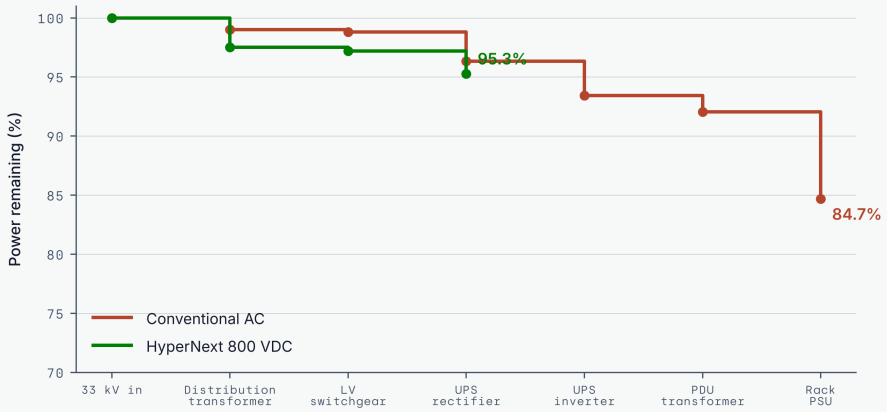


Figure 5. Cumulative power remaining down each chain, plotted from the two tables above. Each step is a conversion and each drop is an asset.

Two caveats belong with these numbers. The first is partial load. Every figure above is a design-point value, and a data hall rarely runs at its design point. Converter efficiency falls away at light loading, often sharply below about a quarter of rating, so the honest comparison between two architectures is made across a load profile rather than at a single point. The second is that these are design values. Measured commissioning figures from HN-1 Phase 1 will be published when the campus is energised, and where they differ from this table we will say so.

The second payment is the one that is usually left out. A watt lost in conversion becomes a watt of heat inside the building, and removing it costs the cooling plant a further fraction of a watt. Against a design power usage effectiveness ceiling of 1.30, every watt of conversion loss carries about 1.3 watts at the meter. That is the mechanism by which an electrical decision shows up in the cooling bill, and it is why the two designs cannot sensibly be optimised apart.

06

Copper, iron and floor area

Three materials arguments follow from the same decision, and all three are ecological as much as commercial.

Copper is sized by current, not by power. A conductor is chosen for the amperes it must carry and the temperature rise that produces, which means the same megawatts delivered at a higher voltage need less metal. The comparison that shows this most starkly is at the rack. A 120 kilowatt rack fed at 800 volts draws about 150 amperes. The same rack fed at 48 volts draws about 2,500. That ratio is why the industry moved to higher rack voltages at all.

The larger saving is structural rather than arithmetic. Taking medium voltage to the rack in a single assembly removes whole cable runs: the low-voltage feeders from the transformer to the switchgear, the switchgear to the uninterruptible supply, the supply to the distribution units. Those runs are not made thinner, they are not installed. HyperNext does not publish a single percentage for copper saved, because it depends on the geometry of each hall and a single figure would be a worse answer than the mechanism.

Copper matters ecologically out of proportion to its tonnage. It is among the most energy and water intensive metals to extract and refine, and availability of it is one of the real constraints on how fast data centre capacity can be built anywhere in the world. Metal not specified is the only metal certain not to be mined.

55% OF GREY SPACE, REMOVED

In a conventional facility, 55 per cent of the grey space is taken by UPS rooms, battery racks, UPS plant and electrical panels. None of it is built here, so none of its floor area is lost. The figure does not include chiller plant.

Floor area is the commercial expression of the same decision. Grey space is built, cooled, insured and maintained, and none of it earns revenue. Space that is never given to a battery room is available as white space instead, which is the part of the building a customer pays for. The iron follows the same logic: a distribution transformer that is not specified is several tonnes of steel and copper that is never wound, never shipped, never filled with oil and never disposed of.

WHAT IS NOT CLAIMED

Lithium has not left the building. Server-internal batteries, the rack backup units that ride out a transfer and ordinary office uninterruptible supplies all still contain lithium cells. What is eliminated is the central battery plant sized for the whole facility load, which is where almost all of the lithium mass in a conventional data centre sits. Spent cells from the units that remain go to authorised recyclers.

07

References and standards

01 **IEEE 1709-2018**
Recommended Practice for 1 kV to 35 kV Medium-Voltage DC Power Systems on Ships. A shipboard document, cited because it is the most mature published practice for medium-voltage direct current. The data centre industry has no equivalent of its own.

02 **IEC 62477-1**
Safety requirements for power electronic converter systems and equipment. Governs the converter assemblies described in sections 2 and 3.

03 **IEC 60664-1**
Insulation coordination for equipment within low-voltage systems. Clearance and creepage at 800 V are set from this.

04 **IEEE 1584-2018**
Guide for Performing Arc-Flash Hazard Calculations. Its method is derived for alternating current and does not transfer directly to a DC network, which is part of the gap described in section 7.

05 **IEC 60617**
Graphical symbols for diagrams. Figure 1 follows this symbol set.

06 **OCP Diablo, v0.5.2, May 2025**
Open Compute Project sidecar power rack, contributed by Google, Meta and Microsoft. Specifies plus or minus 400 VDC nominal with 800 VDC as an alternative, for IT racks from 100 kW to 1 MW. HyperNext campuses standardise on 800 VDC, above the nominal this specification settled on.

07 **ASHRAE TC 9.9**
Thermal guidelines for data processing environments. The ambient basis for the loss and cooling figures in section 5.

08 **HyperNext Research, HN-RP-002**
800VDC Power for the AI Rack Era. The architecture paper this one sits beneath.

09 **HyperNext Research, HN-RP-006**
Liquid Cooling at 600 kW per Rack. Carries the cooling side of the loss argument made in section 5.

HYPERNEXT RESEARCH

Medium-Voltage Solid-State Transformation

Cite as: HyperNext Research, HN-RP-011, "Medium-Voltage Solid-State Transformation", version 1.0, 1 December 2026. 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.