HyperNext Data Centers

RESEARCH PAPER

HN-RP-002

800VDC: Power Architecture for the AI Rack Era

Why every 415 VAC distribution path is now a stranded asset

Conversion arithmetic from grid to GPU. Twenty points of efficiency from architecture alone. What the rack-scale systems of 2027 and 2028 actually require.

This is what is Next.

Series	HyperNext Research
Paper	HN-RP-002
Issued	15 September 2025
Version	1.0
Classification	Public release
Citation	HyperNext Research, HN-RP-002

800VDC: Power Architecture for the AI Rack Era

This paper is part of the HyperNext Research series. Methodology, assumptions, and source data are stated openly so other operators can reproduce the analysis on their own facilities. Citation as "HyperNext Research, HN-RP-002" is welcome.

Contents

1. The rack power budget is not the same problem anymore

ABSTRACT

Design power per rack in hyperscale AI data centres will move from around 22 kW in 2018 to 600 kW by 2027. That is not an incremental change to the existing power architecture. It is a break. This paper walks the energy budget from the high-voltage utility feed to the GPU package, accounts for every conversion loss along the path, and compares the conventional six-stage 415 VAC architecture with the four-stage 800 VDC architecture HyperNext is building. The end-to-end efficiency gap is roughly twenty percentage points. Conventional delivers 75 percent of grid energy to the chip. 800 VDC delivers 95. When inference is billed by the token, twenty points compounds into real money. We cover where the savings come from, what the architecture costs in capital and design effort, and why the rack-scale GPU systems shipping in 2027 and 2028 were architected around DC delivery from the start.

Cooling gets most of the attention in AI data centre discussions. Power architecture is the bigger problem and it sits one layer beneath the cooling problem.

Air cooling stops working as a primary thermal medium around 50 kW per rack. The industry has accepted this. Direct-to-chip liquid cooling is now mandatory for any rack hosting GPUs designed in 2024 or later. The power architecture transition has had less attention but it constrains the cooling architecture in ways the cooling discussion alone does not surface.

Look at the rack power trajectory on a five-year horizon and the constraint becomes obvious.

Year	Reference system	Rack power (kW)	GPUs per rack	Per-GPU power (W)
2018	V100 SXM2, 8-way	22	32	300
2020	A100 SXM4, 8-way	30	32	400
2023	H100 NVL72 (Hopper)	120	72	700
2025	B200 NVL72 (Blackwell)	132	72	1,000
2H 2026	Vera Rubin NVL144 (Oberon)	~200 (est)	144 dies	~1,400
2H 2027	Vera Rubin Ultra NVL576 (Kyber)	600	576 dies	~1,050

Thirty times the rack power over a decade. 415 VAC distribution does not survive that scaling, no matter how it is engineered. The rest of this paper explains why.

The basic problem with AC at rack scale

Conventional data centre power architecture was designed for the 5 to 30 kW rack era. It takes utility voltage (11 kV or 33 kV in India, depending on facility capacity) and steps it down through a sequence of conversion stages. Building transformer to 415 V three-phase AC. Switchgear distribution. UPS layer with battery backup. PDU layer per row. Server PSU inside each chassis producing the 12 V DC the silicon actually wants. Each stage runs at 95 to 98 percent efficiency. The aggregate is the product of all of them.

The conventional path aggregates to 75 to 77 percent. Twenty-five percent of every megawatt arriving at the building edge dissipates as heat before it does any computational work. That dissipated power becomes cooling load. The cooling load needs power to remove. Cooling power adds to total facility power. PUE degrades. The cooling and power architectures are coupled through this loss term in a way that becomes acute at high density. You cannot solve one without addressing the other.

2. The conversion arithmetic

Conventional 415 VAC path

The standard data centre conversion path from 11 kV grid feed to the 12 V DC the silicon wants runs through six stages. Each stage has an efficiency set by physics, by component design, and by operating point. The numbers below are representative best-in-class values for utility-scale equipment in 2025. They are not aspirational.

Stage	Voltage in	Voltage out	Type	Efficiency
1. Building transformer	11 kV	415 V	AC to AC step	98.0%
2. LV switchboard	415 V	415 V	AC distribute	99.0%
3. UPS (online double-conv.)	415 V	415 V	AC->DC->AC	93.5%
4. PDU	415 V	415 V	AC distribute	98.5%
5. Server PSU	415 V	12 V	AC to DC	92.0%
6. Point-of-load regulators	12 V	0.7-1.2 V	DC to DC	88.0%
Aggregate				74.8%

The dominant loss terms are the UPS double conversion (6.5 percent lost) and the server PSU AC to DC stage (8 percent lost). Both stages do the same kind of work twice. The UPS rectifies AC input to DC for the battery interface, then inverts back to AC for distribution. The server PSU then rectifies that AC back to DC for the silicon. Those two terms together cost approximately 14 percentage points of efficiency. That is the headline.

HyperNext 800 VDC path

The 800 VDC architecture removes two of the six stages and operates the remaining stages at higher efficiency. The path collapses to four:

Stage	Voltage in	Voltage out	Type	Efficiency
1. Building transformer	11 kV	415 V	AC to AC step	98.0%
2. Rectifier (with storage)	415 V	800 V	AC to DC	98.5%
3. DC busway	800 V	800 V	DC distribute	99.5%
4. Point-of-rack DC-DC	800 V	48 V	DC to DC	97.0%
5. Point-of-load regulators	48 V	0.7-1.2 V	DC to DC	97.5%
Aggregate				94.2%

Three things change. No UPS double conversion. The rectifier produces 800 V DC directly, the battery storage attaches to that bus through a smaller DC-DC interface, and the bus distributes DC end-to-end. The point-of-load regulators operate on 48 V DC rather than 12 V. That is the OCP rack standard and it gives 5 to 7 percentage points of efficiency improvement at the regulator. Higher input voltage means lower input current, which means lower I squared R losses. Distribution losses on a DC bus at 800 V are negligible compared with three-phase AC at 415 V carrying the same power. Higher voltage, lower current for the same power. DC has no reactive component to compensate.

The 19.4 percentage points of headroom this creates is not marginal. A rack on the 800 VDC architecture draws 600 kW from the grid. The same rack on 415 VAC draws 760 kW. Same computational output. At 1.6 GW of facility capacity, the energy savings work out to roughly the output of a 320 MW power plant.

ENERGY BUDGET HEADLINE

- Conventional 415 VAC delivers 75 paise of every grid rupee to the GPU.
- HyperNext 800 VDC delivers 94 paise of every grid rupee to the GPU.
- The difference is structural. It comes from removing conversion stages, not from squeezing efficiency out of stages that remain.
- For inference workloads billed at INR 0.6 per million output tokens, the recovered nineteen percent generates approximately INR 95 crore per year of additional revenue per gigawatt of facility capacity.

3. What this costs

Capital

The 800 VDC architecture is not free. Per megawatt of IT capacity, capital cost runs roughly 12 to 18 percent higher than 415 VAC. The premium has three sources. The rectifier-storage assembly replaces both the UPS and a portion of the PDU function and costs more per kW than the discrete equivalents. DC switchgear and protection at 800 V costs more than equivalent AC devices because the supply chain for high-voltage DC components in data centre applications is still maturing. The DC busway costs more per metre than its AC equivalent for the same reason.

The premium is shrinking year over year as DC-native components reach scale. Based on supplier roadmaps from Delta, Eaton, and Schneider, we expect it under 8 percent by the end of 2027. At that point the 800 VDC architecture is the lower-cost option without counting any energy savings.

Operating

The nineteen-point efficiency improvement translates to operating cost directly. At INR 8 per kWh for industrial electricity in Telangana and Andhra Pradesh, the saved power on a 1 GW IT-load facility comes to roughly INR 1,330 crore per year of electricity not purchased. Payback on the architectural premium runs 14 to 22 months at current electricity prices.

There is a second-order effect. Power not lost is also heat not generated. Every kilowatt of avoided conversion loss is a kilowatt the cooling system does not have to remove. At gigawatt scale this is not negligible. Roughly 4 to 5 percentage points of additional PUE improvement comes from the lower cooling load. That feeds back into the water consumption discussed in HN-RP-001.

Operations

The 800 VDC architecture is operationally simpler than 415 VAC, not more complex. One bus voltage to manage instead of three. No phase balance issues because there are no phases. No power factor to correct because there is no reactive power. Protection coordination is the hardest part of any data centre electrical design and it gets simpler under DC. Fewer fault types. Well-characterised fault current behaviour.

The trade-off sits in the workforce. Indian electrical engineering training is overwhelmingly AC-centric. DC competence has to be built deliberately, person by person, and most operators have not started. HyperNext runs a partnership with two Indian engineering colleges on a DC-power curriculum. The training materials are available to other operators on request. This is not a competitive position. It is industry infrastructure.

4. Why the rack-scale systems require it

Vera Rubin Ultra NVL576

The NVIDIA Vera Rubin Ultra NVL576 platform, scheduled for the second half of 2027, is the reference AI infrastructure system for that horizon. A rack-scale system: 576 GPU dies in a single liquid-cooled cabinet, drawing 600 kW continuous per the figure Jensen Huang confirmed at GTC 2025. The platform uses 48 V DC at point-of-rack with a high-voltage DC distribution upstream. The Kyber rack infrastructure is built for this density.

The DC-native decision came from NVIDIA running the same arithmetic this paper has run. At 600 kW per rack, the current at 415 VAC three-phase would be approximately 835 amperes per phase. Cable sizing for that current is a 240 square millimetre conductor, three of them, plus neutral and ground. The mechanical envelope of a four-conductor 250 sq mm AC distribution path does not fit the floor space and weight budget of a modern AI rack. At 800 VDC the current is 825 amperes on a single two-conductor path. The mechanical envelope is half the size and a third of the weight.

That is the deeper reason the AC versus DC argument is settled. Not the efficiency, though the efficiency matters. The physical reality. Running 600 kW into a single rack requires either DC at high voltage or AC at voltages no data centre is willing to handle. The industry has chosen DC.

AMD Helios

The AMD Helios rack-scale platform, which HyperNext deploys alongside the NVIDIA Vera Rubin family at Kakinada, takes a different rack form factor than Oberon and Kyber. Helios is a double-wide rack built to the Open Rack Wide standard contributed by Meta, with 72 MI450 series GPUs per rack. The MI450 GPU TBP sits around 1,200 to 1,400 watts. Total rack power is in the 110 to 150 kW range based on financial-analyst modelling and engineering bottom-up calculation. AMD has not yet published an official rack-level power figure. Helios at this density does not need 800 VDC distribution to make the cable arithmetic work, but it does require high-density direct liquid cooling and 48 V DC at point-of-rack. Both NVIDIA and AMD have moved decisively away from 12 V DC at the rack, which is the relevant architectural commonality.

Two independent vendors converged on the same architecture. That is not coincidence. Both companies arrived at the same answer to the same engineering question. For data centre operators the implication is unambiguous. Any facility being built between now and 2030 to host these systems must commit to DC distribution at construction. Retrofitting from AC to DC after the fact is not economic. The choice is made at white-paper stage, not at deployment stage.

5. What we recommend

For new-build AI data centre projects starting design today, the architectural principles below are the ones we would commit to.

1. **Specify 800 VDC as the primary bus voltage at design phase.** Retrofitting later is uneconomic. The decision has to be made before structural design is locked.
2. **Co-locate rectifier and battery storage in a single integrated assembly.** This is what removes the UPS double-conversion penalty. Buying rectifiers and batteries separately and integrating them on-site reintroduces the loss term that 800 VDC was supposed to eliminate.
3. **Plan for 48 V DC point-of-rack delivery.** The OCP rack designs assume this. Designing the white space for 12 V or 24 V rack input reintroduces a regulation stage that does not need to exist.
4. **Buy DC switchgear and protection from established vendors.** The supply chain is mature enough now (Delta, Eaton, Schneider ship at scale) that single-sourcing the critical components from a tier-one OEM is the right call.
5. **Train the operations team on DC.** Operational discipline is different. Fault behaviour is different. Maintenance procedures are different. Treat the architecture transition as a workforce transition.

HEADLINES

> The conversion path from grid to GPU runs six stages in conventional 415 VAC and four stages in the 800 VDC architecture HyperNext is building.

> Aggregate end-to-end efficiency is 75 percent conventional, 94 percent for 800 VDC. The 19-point gap is structural.

> At 1 GW facility scale, the gap is worth roughly INR 1,330 crore per year of electricity not purchased, plus 4 to 5 additional points of PUE improvement from reduced cooling load.

> Rack-scale AI platforms shipping in 2027 and 2028 require 800 VDC input as a hard architectural constraint, not a preference.

> For new-build projects, the architectural decision has to be made at design phase. Retrofitting is uneconomic.

The next paper in the series, HN-RP-003, looks at the third major architecture decision facing Indian data centre operators between now and 2030. Sovereignty. Indian data residency law has tightened substantially in the past three years. The argument for sovereign cloud is increasingly an operational constraint, not a marketing position. HN-RP-003 walks through what that constraint means in practice for BFSI, government, and regulated enterprise workloads.

6. DC fault behaviour and protection

Operators familiar with AC distribution often raise protection coordination as the most worrying part of the transition to DC. The concern is legitimate. DC fault current behaviour is genuinely different from AC and the protection devices that have decades of operational track record on AC do not directly translate. The section below covers what changes and how the HyperNext design handles it.

Why DC faults look different

In an AC system, the fault current waveform passes through zero twice per cycle. This natural zero crossing is what allows AC circuit breakers to interrupt the fault. The breaker mechanism does not have to forcibly drive the current to zero. It just needs to open the contacts during a zero crossing and let the arc extinguish naturally.

DC has no zero crossing. A fault current at 800 V DC will rise to its steady-state value (limited by source impedance and cable impedance) and stay there until something forcibly interrupts it. Conventional AC breakers will not clear a DC fault. The arc continues to draw current across the opening contacts and the breaker melts before the fault is cleared.

The two engineering approaches to DC fault interruption are mechanical hybrid breakers (which use parallel solid-state switching to force a zero crossing) and pure solid-state breakers (which interrupt directly through semiconductor switching). Both are now available from tier-one DC infrastructure suppliers (Delta, Eaton, Schneider). HyperNext uses solid-state breakers at the rectifier output and hybrid breakers at the rack-level DC-DC stage. The selection is based on interruption time, fault current handling, and cost.

Fault current calculation

The peak fault current in an 800 VDC system is limited by three impedances in series: the rectifier output impedance (typically 5 to 15 milliohms for a 1 MW rectifier-storage assembly), the cable impedance from rectifier to fault location (depends on conductor cross-section and length), and the fault impedance itself.

For a representative HyperNext distribution: 800 V source, 0.01 ohm rectifier impedance, 0.005 ohm cable impedance over 30 metres of 95 sq mm conductor, and a bolted fault (zero impedance). Peak fault current is:

```
I_fault = V_source / (R_source + R_cable + R_fault)
       = 800 / (0.010 + 0.005 + 0)
       = 800 / 0.015
       = 53,333 A   (53 kA peak)
```

The protection device at the rectifier output must interrupt this current before thermal damage occurs at any point in the circuit. The HyperNext design uses solid-state breakers rated for 65 kA interrupting capacity at 800 V DC, with interruption time below 100 microseconds. The cable thermal rating for short-circuit conditions (the I squared t value) is 1.2 megaamps squared seconds, which corresponds to approximately 425 ms at 53 kA before insulation damage. The breaker interrupts in under 100 microseconds. The thermal margin is greater than 4,000x.

Protection coordination

The coordination problem under DC is conceptually simpler than under AC because there is no current asymmetry from inductive sources to handle. The arithmetic is straightforward: each downstream device must interrupt faster than its upstream device for any fault current within the downstream device range, so that the upstream device does not need to operate.

The HyperNext design uses a four-level coordination: rectifier output breaker (65 kA, 100 microseconds), busway protection (32 kA, 200 microseconds), rack-input breaker (10 kA, 500 microseconds), point-of-load fuse (1 kA, 2 ms). A fault at the chip level operates only the point-of-load fuse. A fault at the rack input operates the rack breaker. A fault in the busway operates the busway breaker. A fault at the rectifier output is cleared by the rectifier breaker itself. Selectivity is maintained at every level.

Ground fault detection

800 VDC distribution at HyperNext is configured as an ungrounded system. Both poles float relative to chassis ground. This is the standard for high-voltage DC data centre distribution. The advantage is that a single-pole-to-ground fault does not cause immediate fault current flow. The system continues to operate while the fault is located and cleared.

The disadvantage is that ground faults must be detected by some other mechanism. The HyperNext approach uses two independent ground-fault detection methods. Active injection (a low-current AC signal injected to ground and monitored for return) provides continuous monitoring. Insulation resistance monitors at each branch panel detect degradation before a hard fault occurs. Both methods alarm to the BMS at thresholds that allow planned isolation rather than emergency shutdown.

7. Cable sizing and thermal analysis

Cable sizing under DC follows the same fundamental relationships as AC (resistance times current squared equals heat) but the absence of skin effect and proximity effect at DC means the same conductor carries more continuous current than its AC rating implies. The benefit is real but modest. The DC ampacity advantage is typically 5 to 12 percent over the AC rating of the same conductor in the same installation.

Sizing methodology

For a busway carrying continuous DC current, the conductor sizing is determined by the lower of two limits. The voltage drop limit (the conductor cannot drop more than the design tolerance, typically 2 percent end-to-end) and the thermal limit (the conductor temperature cannot exceed the insulation rating under continuous operation).

```
VOLTAGE DROP CALCULATION
  Acceptable drop:    2% of 800 V         =  16 V
  Continuous current: 1000 A
  Path length:        50 m one-way         = 100 m round trip

  Required resistance: R = V/I = 16/1000   = 0.016 ohm
  Resistance per metre: 0.016/100          = 0.16 mohm/m
  Required conductor cross-section: ≥ 100 sq mm copper

THERMAL CALCULATION
  Continuous current: 1000 A
  Conductor: 95 sq mm copper, free-air installation
  Resistance per metre: 0.193 mohm/m at 20°C
  Heat dissipation per metre: I² R = 1000² × 0.000193 = 193 W/m

  Required ampacity for 70°C conductor temperature: ≥ 950 A
  Actual ampacity (95 sq mm, free air, 30°C ambient): 410 A AC, 440 A DC

  Thermal check fails for 95 sq mm. Move to 185 sq mm:
  Ampacity (185 sq mm, free air, 30°C ambient): 700 A AC, 750 A DC
  Still fails. Move to 300 sq mm:
  Ampacity (300 sq mm, free air, 30°C ambient): 950 A AC, 1020 A DC
  Passes.

FINAL SELECTION
  Conductor: 300 sq mm copper
  Bound by: thermal limit, not voltage drop
  Voltage drop margin: 16 V allowed, 9.6 V actual at 1000 A (1.2%)
```

The example shows the common case: voltage drop is easier to satisfy than thermal limit. The 300 sq mm conductor is required for thermal compliance even though a 100 sq mm conductor would handle the

voltage drop. In free-air installation the difference is moderate. In conduit or busway enclosure the thermal limit becomes much more restrictive and conductor sizes go up further.

Why DC busway sizing matters

At the rack-input level, a 600 kW rack at 800 V DC needs 750 amperes of continuous supply. Cable sizing for that current with adequate margin is a 240 sq mm conductor pair plus protective earth. The mechanical envelope is approximately 50 mm by 50 mm cross-section per phase. Compare this to 415 VAC three-phase distribution for the same power: 920 amperes per phase, three phases plus neutral plus earth, 185 sq mm per conductor in three-phase configuration. The total cable cross-section is roughly double the DC equivalent.

The mechanical and thermal implications of this are substantial. DC busway penetrations through fire-rated walls are smaller and easier to seal. DC busway termination labour is simpler. DC busway thermal load on the room ambient is lower. These are not the headline efficiency arguments but they are real construction and operational savings.

8. Worked example: 100 MW DC distribution sizing

The worked example below sizes the complete 800 VDC distribution architecture for a 100 MW IT load. The numbers are representative of one floor of the HyperNext Hyderabad Phase 1 build. The methodology generalises to any scale.

Inputs

Parameter	Value
IT load (continuous, design max)	100 MW
Rack count	167 racks at 600 kW each
White space layout	Hot-aisle containment, racks distributed across rows
Cooling architecture	Direct-to-chip with CDU per row
Bus voltage	800 V DC
Voltage tolerance	+/- 2% at any point
Redundancy	2N at rectifier-storage level, N+1 at rack-DC stage

Rectifier-storage sizing

The rectifier-storage assembly converts incoming 415 VAC to 800 VDC and integrates battery storage for the UPS function. For 2N redundancy on 100 MW IT load:

```
TOTAL CAPACITY REQUIRED
  IT load:                                  100 MW
  Distribution losses (4 to 6%):              5 MW
  Rectifier internal losses (1.5%):         1.6 MW
  Total upstream load:                      106.6 MW

  2N redundancy: two independent rectifier assemblies, each rated for full load
  Each assembly capacity:                   107 MW (rounded)

STORAGE SIZING (UPS RUNTIME)
  Hold-up time required: 12 seconds @ full load (until genset start completes)
  Plus 60 seconds for genset stabilisation:   72 seconds total

  Storage energy:  100 MW × 72 s / 3600 = 2.0 MWh
```

```
  Lithium iron phosphate (LiFePO4) cells, 800 V DC bus voltage
  Per-cell voltage 3.2 V nominal:             250 cells in series
  Pack capacity:                              2.0 MWh total
  Per-pack cells:                             250 series x 32 parallel = 8,000 cells per
  Cell capacity:                              80 Ah at nominal voltage

  Number of storage packs (2N):               2 packs, each 1.0 MWh
```

Busway sizing

Each rectifier-storage assembly feeds a main busway that runs the length of the row aisle. For 100 MW IT load split across 15 rows, each row busway carries approximately 6.7 MW. At 800 V DC that is 8,375 amperes per row busway.

```
ROW BUSWAY SIZING
  Continuous current per row:    8,375 A
  Voltage drop budget:           1% over 60 m row length
  Required conductor area:       Approximately 1,200 sq mm per pole

  Implementation: copper laminated busway, 1,500 A per laminate, 6 laminates
                  per pole. Internal thermal management via heat sinking
                  along the busway length.

  Mechanical envelope:           120 mm wide x 50 mm tall per pole
  Total assembly height:         310 mm including enclosure and PE conductor
  Weight:                        38 kg per metre

RACK DROP SIZING
  Per-rack current:              750 A at 800 V DC (600 kW)
  Drop length:                   4 m from busway to rack input
  Conductor:                     240 sq mm copper cable
  Termination:                   solid-state breaker at busway tap-off (10 kA breaking, 2
```

Total architecture summary

Component	Quantity	Unit rating	Notes
415 V to 800 V rectifier-storage	2	107 MW each	2N redundant
Storage capacity	2 packs	1.0 MWh each	72-second hold-up
Main busway (rectifier to floor)	2	120 kA continuous	One per rectifier
Row busway	15	9 kA continuous	One per row

Rack drop cable	167	900 A continuous	240 sq mm
Solid-state breakers (rack tap-off)	167	10 kA breaking	200 microsecond interruption
Point-of-rack DC-DC (800 V to 48 V)	167	600 kW each	N+1 internal redundancy per rack

The capital cost estimate for this architecture is INR 240 crore for the complete distribution (rectifier-storage, busway, rack drops, breakers, point-of-rack DC-DC). Compared with the equivalent 415 VAC architecture (UPS lineup, transformers, switchgear, PDUs), the premium is approximately 14 percent of total cost. At HyperNext electricity prices, payback on that premium runs 19 months.

9. Standards landscape and references

The 800 VDC architecture for data centres is no longer a frontier specification. The standards landscape has caught up substantially over the past five years. The references below are the documents on which the HyperNext design relies.

DC distribution standards

- **IEC 62933 series.** Electrical energy storage systems. Provides the framework for the rectifier-storage integration including battery interface, fault behaviour, and protective device coordination.
- **IEC 60364-8-2.** Low-voltage electrical installations, Part 8-2. Prosumer electrical installations. Covers DC distribution in commercial premises.
- **IEEE 946.** Recommended Practice for the Design of DC Power Systems for Stationary Applications. The reference for fault current calculation and protection coordination at high-voltage DC.
- **OCP DC Power Specification.** Open Compute Project, current version 2.0. Specifies 48 V DC at the rack input with 12 V or lower at the chip. The HyperNext architecture extends this with 800 V DC at the floor level.
- **ETSI EN 300 132-3-1.** Environmental Engineering. Power supply interface at the input to telecommunication and datacom equipment. Up to 400 V DC. The European telecoms DC standard whose 380 V variant influenced the global rack-level DC ecosystem.
- **EMerge Alliance Data/Telecom Center 380 V DC standard.** The 380 V DC predecessor to the 800 V architecture for data centre applications.

Component vendor specifications

The efficiency numbers and protection device characteristics cited in this paper are drawn from publicly available vendor documentation for the following product families.

- **Delta DPR 800 V DC rectifier-storage platform.** Public specification. 95-98% efficiency depending on operating point. UL 1741, IEC 62040 listed.
- **Eaton XR series solid-state breakers.** 800 V DC, up to 65 kA interrupting, sub-100 microsecond interruption time. UL 489 supplemented for DC operation.
- **Schneider DC busway systems.** Up to 12 kA continuous, multiple cross-sections, modular installation. ETIM 8 classified.
- **Vertiv DC switchgear lineup.** Used at the rectifier output stage in the HyperNext design.

NVIDIA and AMD platform specifications

- **NVIDIA Vera Rubin and Vera Rubin Ultra rack-scale platform brief.** Public technical brief, published 2025. The 800 V DC input specification and 48 V DC point-of-rack secondary.
- **AMD Helios rack-scale platform datasheet.** Public datasheet, 2025. The 800 V DC input specification and the rack-level power architecture.

Other sources

- **Uptime Institute.** Data Center Power Efficiency Reports, annual editions 2020 to 2024. The reference for industry-average conversion path efficiency.
- **The Green Grid.** PUE: A Comprehensive Examination of the Metric. White paper WP#49, 2014. The metric definition.
- **Lawrence Berkeley National Laboratory.** Data Center Energy Use: Trends and Drivers, 2021 update. The basis for industry baseline efficiency claims.

HyperNext Data Centers

HyperNext Research

We publish engineering and policy papers because the Indian conversation about AI infrastructure needs more substance than marketing material provides. The papers state methodology openly so other operators can run the same analysis on their own facilities. They report findings that may not flatter the HyperNext commercial position. They get review from the engineering team and the communications partners.

Correspondence on methods, figures, and conclusions: **hello@hypernxt.com**. We read every email.

HN-RP-002 · 800VDC: Power Architecture for the AI Rack Era
15 September 2025 · v1.0

www.hypernxt.com/research
hello@hypernxt.com · +91 99784 23333