HyperNext Data Centers | Research

HN-RP-012

Breaking Direct Current at the Rack

Solid-state protection and the panel that carries it

Why interruption without a current zero is a different problem. Selectivity when every device is fast. What removing the arc does to maintenance.

SERIES	HyperNext Research
PAPER	HN-RP-012
ISSUED	16 December 2026
VERSION	1.0
CLASSIFICATION	Public release
CITATION	HyperNext Research, HN-RP-012

This is what's Next.®

Breaking Direct Current at the Rack

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-012" is welcome.

CONTENTS

01 Where the panel sits
02 Interruption without a current zero
03 The breaker as a device
04 Selectivity when everything is fast
05 Arc flash, and what it does to maintenance
06 What else the panel carries
07 References and standards

01 Where the panel sits

The rack power panel is the last piece of fixed electrical infrastructure before a rack. It takes 800 volts direct current from the busway, protects it, meters it, splits it into feeders and holds the battery drawers that ride out a disturbance.

It is not an AI-only product, and that is worth saying because the architecture is usually discussed as though it were. The same panel feeds general purpose servers and storage. On a HyperNext campus a colocation rack typically draws between 12 and 18 kilowatts and some mainframe-class racks draw under seven, while an accelerated compute rack on the same floor may draw fifty times the latter figure. One panel design has to serve all of them.

Two details in that arrangement carry most of the commercial argument. The output cassettes plug in and are hot-swappable, so the current rating presented to a rack is changed by exchanging a cassette rather than by replacing a panel, a board and a breaker in a live hall. And the frame is sized ahead of its present duty: a single 500 ampere feeder is installed today against a frame built to accept a second. Capacity that is not yet needed does not have to be bought, but the room for it has to be designed in from the start, because the frame is the part nobody wants to change later.

FIGURE 1

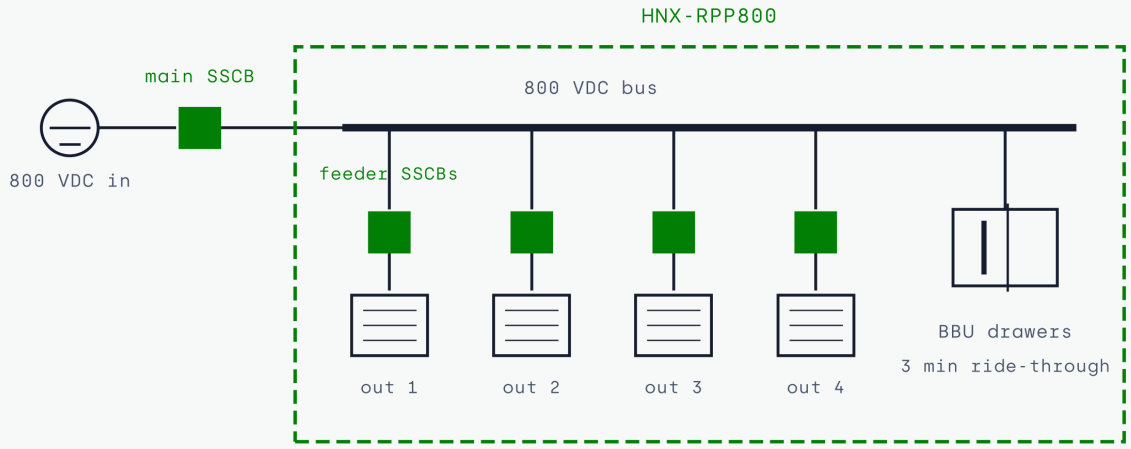


Figure 1. Internal arrangement of the panel. A main solid-state breaker, an internal 800 volt bus, feeder breakers ahead of each output cassette, and the battery drawers inside the same enclosure.

WHERE THIS PANEL DOES NOT GO

Rack-scale accelerated systems from the major vendors arrive with their own power shelves and internal distribution already built in. The HyperNext panel is not placed in front of those. It serves the rest of the floor, which on any real campus is most of the floor.

02 Interruption without a current zero

An alternating current circuit crosses zero one hundred times a second. Every mechanical breaker ever built for alternating current exploits that fact, whether its designers thought about it in those terms or not.

The sequence is familiar. Contacts part, an arc strikes across the gap, and the arc persists until the current reaches its natural zero. At that instant there is no current to sustain ionisation, the arc extinguishes, and the gap recovers dielectric strength before the voltage rises again on the next half cycle. Clearing takes tens of milliseconds and the device can be comparatively simple, because the waveform is doing the difficult part.

Direct current offers no such instant. Current has to be forced to zero,

or the arc has to be stretched, cooled and lengthened until the voltage available can no longer sustain it. Mechanical direct current breakers do exist and they work, using commutation circuits that inject a counter-current to create an artificial zero. They are larger, slower and more expensive than their alternating current equivalents, and at 800 volts in a dense hall they are the wrong answer.

A third option exists and deserves a fair hearing, because it is what most of the industry reaches for first. A hybrid breaker puts a mechanical contact and a semiconductor branch in parallel. The contact carries the current normally, so the standing loss is almost nothing, and on a fault the current is commutated into the semiconductor branch while the contact opens into near-zero current. It is an elegant arrangement and it genuinely works.

It is also slower than it looks. The commutation itself is fast, but the mechanical contact still has to travel, and the device cannot be declared open until it has. That puts clearing back into the hundreds of microseconds at best, and it keeps a moving part in the protection path with the wear and the maintenance interval that implies. For a campus bus where the economics turn on how little is lost during normal running, the hybrid is the right answer. For a feeder to a rack, where the current is a few hundred amperes and the thing being protected costs more than the breaker by two orders of magnitude, it is not.

A solid-state breaker takes a different route entirely. A semiconductor switch in the conducting path is commanded off. There are no contacts, so there is no arc to extinguish, and the transition takes microseconds rather than milliseconds.

FIGURE 2

The dashed line is what the circuit would deliver if nothing interrupted it.

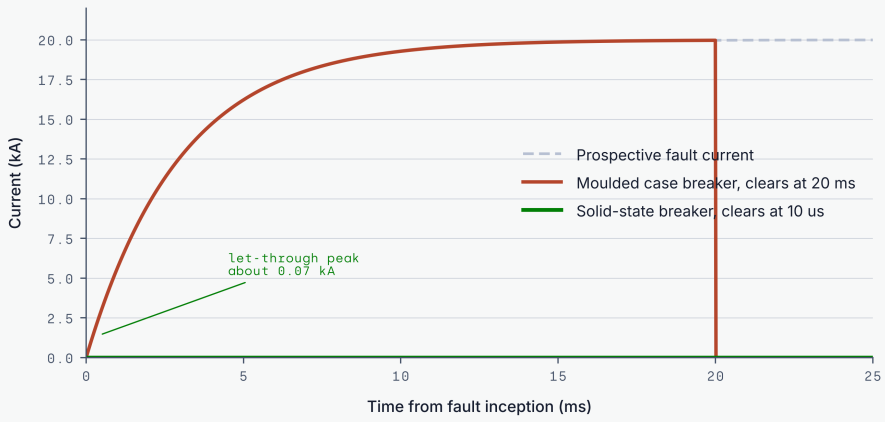


Figure 2. Current against time from fault inception, at a 20 kA prospective fault. The limiting device never allows the current to develop, which is why the comparison is better made on let-through energy than on peak current.

3 orders LESS LET-THROUGH ENERGY

Damage at a fault scales with the square of current and with the time it is allowed to flow. Clearing in ten microseconds rather than twenty milliseconds is the difference between a replaced module and a replaced lineup.

03

The breaker as a device

A solid-state breaker is a power semiconductor, a gate drive that can turn it off into a fault, a means of sensing that fault fast enough to be useful, and a way of absorbing the energy that has nowhere else to go.

FIGURE 3

Rendered from the issued manufacturing model.

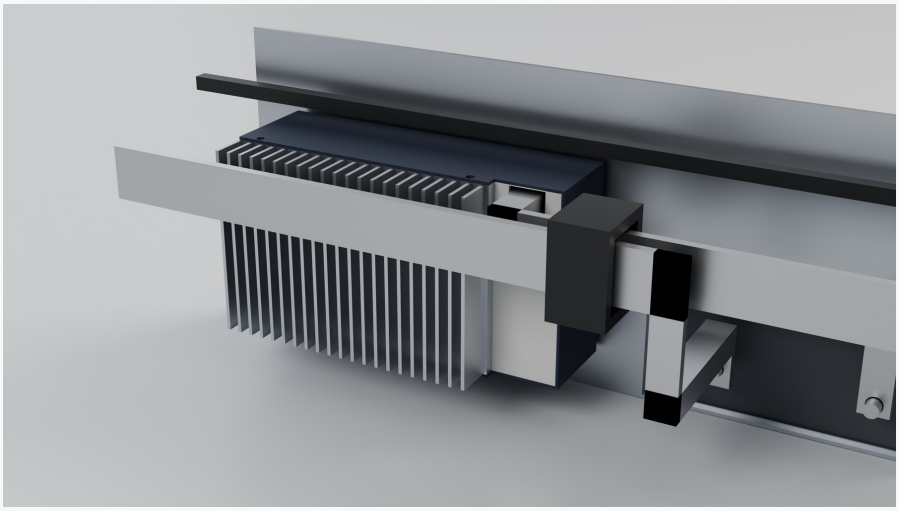


Figure 3. The solid-state breaker assembly as designed for the HyperNext panel bay.

The device itself is the least interesting part. Wide bandgap semiconductors give the blocking voltage and the switching speed, and the engineering sits around them. Sensing has to resolve a developing fault in a few microseconds, which rules out any measurement with a filter long enough to be comfortable. The gate drive has to turn a device off while it is carrying many times its rated current, which is the opposite of the condition it is usually characterised for.

Sensing deserves more attention than it usually gets. A shunt resistor is accurate and cheap and sits in the current path, where its own inductance becomes part of what it is trying to measure at the frequencies that matter here. A Hall effect or fluxgate sensor is isolated and introduces no loss, and has a bandwidth limit that has to be

checked against the fault it is meant to catch rather than against the steady current it is meant to meter. In practice the panel does both, because the measurement that bills a customer and the measurement that trips a breaker have almost nothing in common beyond the quantity they share.

THE ENERGY HAS TO GO SOMEWHERE

Interrupting a current in microseconds does not remove the energy stored in the inductance of the cable run feeding the fault. That energy appears as a voltage across the opening device, and if it is not clamped it will destroy the thing that just did its job. Snubbing and surge arrest across the breaker are not accessories. A specification that quotes interruption time without stating clamping energy is incomplete.

There is also a standing cost. A mechanical breaker that is closed is a piece of metal, and its on-state loss is close to nothing. A semiconductor that is on has a forward voltage drop and dissipates power for every hour the facility runs. That loss is small in percentage terms and it is real, it has to be cooled, and it is part of why the panel has a designed air path rather than an incidental one.

FIGURE 4

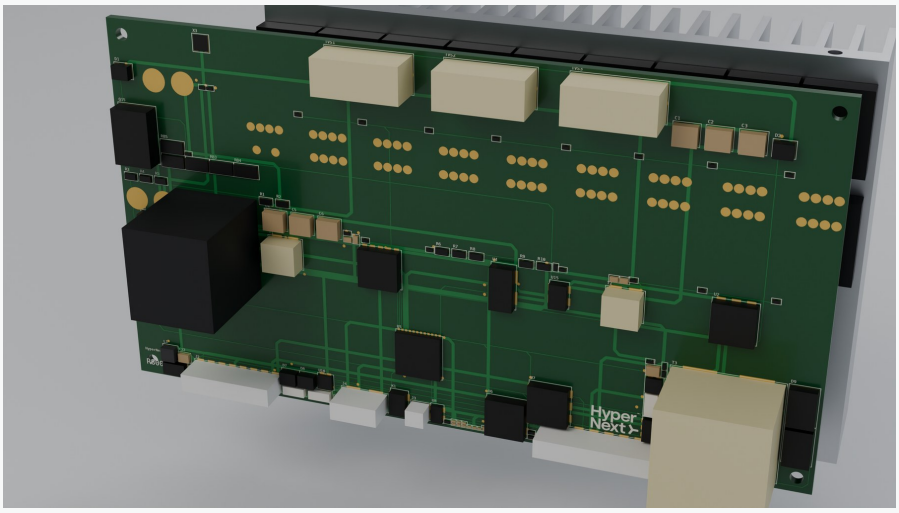


Figure 4. The breaker control board. Sensing, gate drive and the protection logic sit together so the decision path is as short as the physics requires.

04

Selectivity when everything is fast

Selectivity is the property that a fault is cleared by the device closest to it, and not by its parent. In an alternating current network this is arranged on time. In this network there is no time to arrange it with.

The conventional method grades devices by their time-current curves. A downstream breaker is set to operate faster than the one above it at any given current, and the margin between them is measured in tens of milliseconds. That margin exists because mechanical devices are slow, which here is a feature rather than a defect.

When every device in the chain clears in microseconds, the grading margin collapses to nothing. Two breakers seeing the same fault will both decide to operate, and the facility loses a whole feeder to a fault that belonged to one rack. Coordination has to be arranged on something other than time.

- Current threshold. The feeder device is set to a level the downstream device cannot reach, which works when the downstream ratings are well separated and works poorly when they are not.
- Rate of rise. The gradient of current tells a device something about where the fault is, because inductance between the source and the fault limits how fast current can develop. This is information a mechanical device never had time to use.
- Explicit communication. Devices tell each other what they have seen. This is the most deterministic approach and it introduces a dependency on a communication path, which then has to be as reliable as the protection it serves.

The rate-of-rise argument is worth unpacking, because it is the one that does real work. When a fault appears, the rate at which current climbs is set by the driving voltage divided by the inductance between the source and the fault. A fault at the end of a long feeder run sees more inductance than one on the panel bus, so its current climbs more slowly. A device that measures the gradient as well as the magnitude therefore learns something about distance, and distance is exactly what selectivity needs to know. A mechanical breaker never used this information because by the time it could have acted on it the gradient had long since stopped mattering.

The limitation is equally worth stating. Inductance between source and fault is not a number the designer controls to a tight tolerance. Cable routing changes it, busway geometry changes it, and a fault halfway

along a run presents a value between the two cases the analysis was built on. The discrimination this gives is real and it is not sharp, which is why it is used alongside a threshold rather than instead of one.

HyperNext uses the first two and treats the third as a supervisory layer rather than as the primary mechanism, because protection that depends on a message arriving is protection with a new failure mode. The margins involved are narrow enough that they have to be demonstrated in simulation across the full range of fault positions before any of it is energised. Section 7 says what we had to build to do that.

05

Arc flash, and what it does to maintenance

The strongest argument for this architecture is not efficiency and not even availability. It is what happens to the person who has to work on it.

Arc flash incident energy is, to a first approximation, proportional to how long the arc burns. That is why every mitigation technique in conventional practice is ultimately about reducing arcing time: maintenance switches that lower the instantaneous pickup, zone interlocking, optical sensing that bypasses the normal measurement delay. All of them are attempts to shorten an interval that a mechanical device cannot shorten below its own operating time.

FIGURE 5

Thresholds from NFPA 70E. Plotted on the AC method, for comparison only.

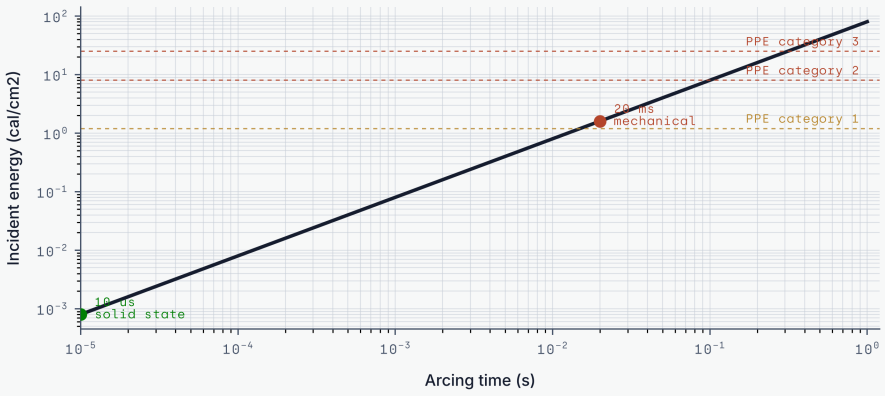


Figure 5. Incident energy against arcing time at fixed arcing current, with the protective clothing category thresholds marked. The relationship is proportional, so the horizontal axis is the whole story.

A fault that is interrupted before an arc can establish does not produce an arc flash in the conventional sense at all. The practical consequence is a change in working method: what protective equipment is required, whether a task can be done energised, how long a permit takes, how

many people are needed. Those are the things that decide whether maintenance happens on schedule.

SAID CAREFULLY

The incident energy method in IEEE 1584 is derived from alternating current testing and does not transfer directly to a direct current network. Figure 3 is plotted on that method to show the shape of the relationship, not to assign a category to this equipment. Direct current arc-flash assessment uses different models, the published test basis behind them is thinner, and HyperNext treats that gap as an open item rather than a solved one.

06

What else the panel carries

Protection is the part that justifies the panel. Most of what it earns its place doing is quieter than that.

TABLE 1

FUNCTION	HOW IT IS DONE	WHY IT IS IN THE PANEL
Metering	Per output, to rack level	Billing and capacity planning need the measurement at the boundary
Fault isolation	Per feeder breaker	A fault takes one rack, not a row
Redundancy	Dual feed into one frame	A source can be lost without a transfer
Ride-through	Battery drawers in frame	Three minutes at rack level
Density change	**Cassette exchange**	**No downstream rework, no downtime**
Monitoring	Integrated to the BMS	One system sees power, cooling and security

What the panel does besides interrupting faults. The cassette row is the one customers notice, because it is the one that removes a negotiation from the lease.

The density argument is worth stating plainly. In a conventional facility a customer who signed for ten kilowatts and now wants thirty is asking for a new panel, a new tap-off box and a new breaker, installed in a live hall, with the downtime that implies. Here the power already sits on the bus. The cassette is exchanged and the rack draws more. The limit is the power envelope of the hall rather than the copper that was installed on the day it was built.

That matters more each year, because the driver is not customers growing their business. It is servers getting denser. The same workload in the same floor area draws more every refresh cycle, and a facility whose distribution was fixed at build time has to be rebuilt around

tenants who have not changed what they asked for.

Metering is the other function that pays for itself quietly. Measuring at the panel output rather than at the row gives a per-rack figure without a separate metering strip inside the rack, which removes a device, a mounting position and a failure mode from every cabinet on the floor. It also means the billing boundary and the protection boundary are the same boundary, so a dispute about consumption and an investigation into a trip are looking at one record rather than two that disagree by a small amount for reasons nobody can reconstruct six weeks later.

All of it reports into the building management system alongside cooling and security. That is an unglamorous integration and it is the one that decides whether an operations team can answer a question at two in the morning. A rack that tripped, a loop whose supply temperature drifted and a door that was opened are three facts that mean very little apart and quite a lot together.

07 References and standards

01 **IEEE 1709-2018**
Recommended Practice for 1 kV to 35 kV Medium-Voltage DC Power Systems on Ships. The nearest mature published practice for direct current protection.

02 **IEEE 1584-2018**
Guide for Performing Arc-Flash Hazard Calculations. Its empirical model is derived from alternating current testing and does not transfer directly to direct current, which section 5 treats carefully.

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

04 **IEC 60947-2**
Low-voltage switchgear and controlgear, circuit breakers. Written around mechanical interruption, and the reference point this design departs from.

05 **NFPA 70E**
Standard for Electrical Safety in the Workplace. The source of the incident energy thresholds plotted in Figure 3.

06 **OCP Diablo, v0.5.2, May 2025**
Open Compute Project sidecar power rack. The HyperNext panel follows the Open Compute mechanical and management conventions while standardising on 800 VDC.

07 **HyperNext Research, HN-RP-011**
Medium-Voltage Solid-State Transformation. Describes what feeds this panel.

HYPERNEXT RESEARCH

Breaking Direct Current at the Rack

Cite as: HyperNext Research, HN-RP-012, "Breaking Direct Current at the Rack", version 1.0, 16 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.