

The Last Few Centimeters

Rewiring AI Infrastructure for the Next Generation of Compute

What Cellink Testing Shows

~6x

lower loop inductance than conventional wire cables, with geometry fixed at manufacture.

~1%

ohmic losses of delivered power, compared with 3.28% for a PCB solution at the same 4.5mm height.

275W

losses saved per 12kW tray compared to a PCB solution.



Cellink's incredibly thin flex harness efficiently carries high currents.

Executive Summary

Every AI factory pays a tax that never appears on a spec sheet. It accumulates around the clock, on every GPU, in the last few centimeters between the power converter and the silicon. Energy lost in the power delivery network becomes heat that must be removed, while voltage margin added to protect against droop draws even more power. At gigawatt scale, these losses can add up to millions of dollars a year, consuming capacity that a power-capped facility could otherwise devote to compute. With the next generation of compute pushing power density higher than ever, the infrastructure behind it will have to evolve.

Fundamentally, AI power inefficiency is a wiring problem, and it's already been solved once. Cellink has seen this constraint firsthand. For more than a decade, the company has built flat, flexible power circuits for electric vehicles, where engineers compete for every millimeter of space and fraction of a percent in efficiency. More than four million Cellink parts are in vehicles today, with a success rate exceeding 99.999%. The same manufacturing technology now underpins its data center portfolio, including PowerPlane, an ultra-thin laminated power delivery network that moves high current beneath the motherboard while reducing electrical losses and freeing space around the GPU.

Electrical performance, however, is only one piece of the equation. Integrating a more efficient power delivery architecture asks platform teams to consider design cycles, qualification requirements, and supply resilience. At the same time, AI infrastructure is already being re-architected around higher power densities, liquid cooling, and rack-scale design, facing many of the same constraints the automotive industry has spent years solving. That convergence makes Cellink's automotive experience especially relevant as the final few centimeters of power delivery become a growing constraint on AI infrastructure.

The Hidden Cable Tax

Power follows a predictable route to a GPU. It enters the rack at high voltage and steps down through a chain of conversions, from 800V to 50V, then to 12V or 6V, and finally to below 1V at the silicon. As voltage falls, current rises sharply. Because resistive losses scale with the square of current, the final few centimeters become especially demanding, carrying thousands of amps at very low voltage.

For years, the industry has turned to utility interconnects, medium-voltage gear, power supplies, and rack-level distribution to optimize stages of the chain. The final leg received less attention, despite introducing its own electrical and physical constraints. Cables, rigid busbars, and heavy-copper PCB layers must carry enormous current into the GPU while minimizing losses across an extremely limited amount of space. Unmitigated, these conditions create three recurring costs:

- 1. Voltage droop.** When GPU load changes rapidly, inductance in the power path limits transient response and voltage can dip. Engineers typically compensate by setting nominal voltage high enough to remain above the processor's minimum operating threshold during these events. CelLink estimates that margin can add 5-15W of draw per GPU.
- 2. Thermal overhead.** Power lost in the delivery network turns into heat that the cooling system must remove. Power usage effectiveness (PUE) accounts for the additional facility energy required to support IT equipment, including cooling. At a PUE of 1.2 to 1.5, each watt lost translates to 1.2 to 1.5 watts of facility power.
- 3. Installation variability.** A cable's electrical behavior can also change with routing. Moving a conductor a few millimeters during assembly alters the geometry of the current loop and therefore its inductance. CelLink measurements show that uncontrolled routing can add up to 20 nH of inductance.

This hidden tax grows with increasing compute density. GPU and platform designers pack compute closer together to preserve high-speed signal integrity, leaving less room for power delivery. As a result, cables, optics, and liquid cooling must increasingly compete for the same limited space around the GPU.

Trained by the Toughest Box on Earth

CelLink solved a closely related geometry problem in another high-current, space-constrained system. The company's flat-conductor technology originated in solar, where engineers needed to move wiring across panels without sacrificing current-carrying capacity. Conventional products couldn't provide the required current in a sufficiently flat form factor, leading CelLink to develop a manufacturing process for thin, flexible conductors.

That technology found a new application when automotive engineers asked whether the same approach could connect battery cells. This question moved CelLink into an even more demanding environment: the EV battery pack, where power delivery must carry high current through extremely limited space while minimizing electrical losses and heat. Every millimeter competes with other components, and each additional watt of loss affects efficiency. Conductors also have to be manufactured at scale, assembled repeatably, and remain reliable for years.

CelLink designed around these constraints through its automotive work while building the manufacturing capabilities required for commercial scale in the process. Flat conductors with defined geometry lend themselves to repeatable placement and automated assembly, enabling high-volume production while meeting the extremely low field failure rates required in automotive applications. The box has changed, but the geometry challenge is identical: moving more current through less space without sacrificing electrical or thermal performance.

The box has changed, but the geometry challenge is identical: moving more current through less space.




Figure 1. At less than 1mm, CelLink flex harnesses move high current through a fraction of the space of conventional cabling.

Why Flat Wins

Five physical characteristics give flat power-delivery architectures an advantage in high-current, space-constrained environments like AI infrastructure:

1. More conductor fits within the same vertical space.

Compared with a standard PCB, PowerPlane combines thicker conductor layers with thinner insulation. For a total available thickness of 4.5mm, PowerPlane has the ability to incorporate eight layers of 500µm copper (4mm of total copper thickness), whereas a traditional PCB would only be able to incorporate 26 layers of 70µm copper (1.8mm of total copper thickness). This provides greater conductive cross-section within the same z-height.

2. Thin conductors improve high-frequency current distribution.

As frequency increases, current density shifts toward the conductor surface, a phenomenon known as skin effect. A wide, thin conductor provides greater surface area relative to its volume than a thick busbar, allowing more of the conductor cross-section to carry current effectively at higher frequencies.

3. Controlled geometry reduces inductance.

Loop inductance depends in part on the spacing between supply and return conductors. In a wire cable, that can change with bends, routing, and installation. A flat laminate fixes the geometry during manufacture, which can lower inductance while giving engineers a more consistent value to design around.

4. Flat surfaces create a direct thermal interface.

Unlike bundles of round conductors, a broad planar surface can mate directly with a cold plate, enabling tighter integration between power delivery and cooling.

5. Flat conductors simplify automated assembly.

Conventional wire harnesses can move, twist, and tangle during handling, introducing variability during routing. Flat parts maintain a defined shape that can support robotic pick-and-place, reducing reliance on manual routing and simplifying placement. Consolidating multiple conductors and connectors into a single flat assembly reduces the number of individual routing and connection steps required so there are few opportunities for error.

The Proof

The physical advantages of flat conductors are reflected in two separate experiments examining different aspects of performance. The first compares the inductance of conventional wire cables with CellLink's narrow flex harness, while the second evaluates electrical losses and thermal performance of a larger, square-shaped flex design. Collectively, these findings demonstrate clear benefits of flexible conductors across multiple configurations:

Lower Inductance

CellLink manufactures its laminate with the supply and return conductors held in a defined position relative to one another, reducing one source of assembly-driven electrical variation. In testing, its flex harness demonstrated substantially lower loop inductance than conventional wire cables across frequencies, measuring 15.8 nH at 100 kHz compared with 80.7 to 103.4 nH for the cable configurations. That represents roughly 5.1x to 6.5x lower inductance.

Same 4.5mm z-height · 2000A × 6V (12kW) load




Figure 2. PowerPlane saves roughly 275W per tray versus a 26-layer PCB power delivery network at the same 4.5mm z-height and 12kW load.

Lower Electrical Losses

PowerPlane architecture packages four independent 3kW electrical paths beneath the motherboard, connecting the DC/DC converter to the GPUs. CellLink compared this architecture with a conventional PCB power delivery network under the same electrical load and 4.5mm z-height. For a 12kW power delivery network, the performance differences are substantial.

Same 4.5mm z-height. Same 12kW load. About 70% less loss.

PowerPlane loss calculations account for 48 connectors, while the PCB model incorporates load-side and PCB-side interfaces. Under these conditions, PowerPlane reduced modeled ohmic losses by about 70%, from 397.5W to 121W. **This translates to a savings of about 275W per GPU tray.**

Lower Thermal Load

Reducing resistance has an additional benefit beyond electrical efficiency: less heat to manage around the GPU.

CellLink’s thermal analysis demonstrates the magnitude of this difference, combining measured PowerPlane performance with calculated results for the PCB solution. Peak temperature rise over ambient reached +63°C for PowerPlane, compared with a calculated +193°C for the PCB design. The 130°C difference illustrates how electrical losses translate into thermal load.

CellLink saw a similar relationship in a separate power-supply test comparing a 3.5mm flex harness with a 5mm PCB busbar. At 35A, the flex harness reached about 29°C, compared with 37.6°C for the PCB busbar. At 40A, temperatures reached 33.2°C and 38.4°C, respectively.

The flex harness achieved these lower temperatures while requiring less vertical space, reducing package height by 30% in an environment where every millimeter around the compute platform is contested.

What the Results Mean for Next-Generation AI Infrastructure

CellLink’s findings demonstrate how rethinking power delivery can improve electrical efficiency, reduce thermal load, and preserve valuable space within increasingly dense AI servers.

By addressing these constraints together, PowerPlane offers a practical alternative to conventional PCB architectures as GPU power requirements continue to rise.

The implications extend beyond individual components. As AI infrastructure scales, power delivery must evolve alongside compute density, supporting higher electrical loads without introducing unnecessary losses, heat, or packaging constraints.

From Nanohenries to Dollars

Electrical losses measured in nanohenries, millivolts, and watts can look small at the component level, but multiplied across thousands of trays, they become a facility-level constraint. The cost accumulates in two ways:

- Higher inductance can increase voltage droop during GPU load transients, leading engineers to add nominal-voltage margin to keep the processor above its minimum operating threshold.
- Resistive losses create a continuous draw of their own, with the resulting heat adding to the facility’s cooling burden.

The PowerPlane comparison provides one way to quantify the impact using ohmic losses alone. CellLink modeled a 276.5W reduction per 12kW tray compared with the PCB solution. At a PUE of 1.25, this translates to approximately:

What 275W per tray is worth at a PUE of 1.25				
346W	3,000 kWh	\$240	\$2.4M	~6 mo.
facility power saved per tray	annual energy savings per tray	per tray each year at \$0.08/kWh	annually across a 10,000-tray deployment	estimated payback, assuming an added \$120 per tray

This excludes potential savings from reduced voltage margin, but the value of those watts isn’t limited to the electricity bill. In a power-capped AI factory, efficiency can translate directly into increased compute headroom.

Supply Resilience is Part of the Architecture

Differentiated power delivery technology creates a legitimate sourcing consideration. AI infrastructure buyers need confidence that critical components will be available across the platform lifecycle, production can scale with deployment needs, and disruption at a single supplier won’t jeopardize system availability.

Commodity components do not guarantee consistent system performance

Standardization and broad availability can reduce sourcing risk, but the resulting power delivery system may still be highly specific to the platform. A 26-layer heavy-copper PCB designed for a particular GPU tray, for example, is already tied to its electrical, mechanical, and thermal requirements. Conventional cables may similarly be available from multiple suppliers, but their electrical behavior can also vary with conductor geometry, routing, and installation. The distinction, therefore, is not simply between proprietary and commodity components. Buyers also have to consider how consistently each architecture delivers the required performance once integrated into the system.

Differentiated power delivery can work within established architectures

CellLink's differentiation lies primarily in how its flat conductors are manufactured and configured. Its products are designed around standard rack mechanics, OCP practices, and existing interfaces, and can use aluminum or copper conductors. This gives engineering teams flexibility in balancing conductivity, weight, cost, and packaging requirements.

Automotive production provides evidence of manufacturing maturity

CellLink's automotive history demonstrates its ability to manufacture high-current conductors at volume and under stringent quality requirements. More than four million CellLink parts are in vehicles today, with a success rate exceeding 99.999%. That manufacturing maturity provides a proven production foundation for deploying differentiated power delivery at scale.

Integrated Power Delivery is the Future

PowerPlane is designed to enable backside power delivery, moving the power path beneath the motherboard and leaving more of the top-side perimeter available for signals, optics, and other critical components. Moving power beneath the board also creates opportunities for closer integration with thermal management. CellLink sees the potential to bring power and cooling to the back of the board together, allowing electrical and thermal interfaces to be considered as part of the same mechanical architecture.

The implications extend beyond PowerPlane. CellLink's manufacturing process can support multiple power delivery formats across the data center, from high-voltage cabling and rack-level distribution to the final connection between power conversion and compute. As AI infrastructure evolves, those applications may increasingly be designed as parts of a coordinated power architecture rather than independent components.




Figure 3. Shown at Open Compute Project Global Summit 2026, CellLink PowerPlane concept enables vertical power delivery for architectures beyond 12kW.

Rewiring AI infrastructure, therefore, isn't simply about moving electricity differently, but designing the physical infrastructure around the demands of the compute it has to support.

Put the Architecture to the Test

The strongest way to evaluate a new power delivery architecture is by testing it against the electrical, thermal, and mechanical requirements of the platform it would ultimately support. With next-generation AI systems already in development, engineering teams have an opportunity to adopt PowerPlane early enough to understand its potential impact on board layout, cooling, available space, and overall system efficiency.

CellLink can work with platform teams to evaluate the technology on their own hardware through joint inductance and thermal testing, evaluation kits, or limited pilots. This gives engineers a way to compare PowerPlane using the metrics that matter to their existing architecture, while identifying integration requirements before the platform design is fixed. To get started, email contact@cllnk.com.

Learn More

Ready to evaluate PowerPlane on your platform? Let's connect.

contact@cllnk.com

cellinktechnologies.com

About CellLink

CellLink develops flat, flexible power circuits for high-current, space-constrained systems. The company's manufacturing technology originated in solar before scaling through high-volume electric-vehicle programs, with more than four million CellLink parts in vehicles today. CellLink is now applying that experience to AI infrastructure, developing innovative power delivery solutions spanning the server tray, rack-to-rack connections, and power shelf to dramatically improve electrical and thermal efficiency.



Evaluate PowerPlane on your own hardware.

Joint inductance and thermal testing, evaluation kits, or limited pilots.

CONTACT

contact@cllnk.com

LEARN MORE

cellinktechnologies.com