Standex EDGE Electronics

INNOVATE / CONSULT / ENGINEER / DELIVER

Reducing Data Center Energy Waste

How custom magnetic design is maximizing Data Center power efficiency and why the move to 48 V and 800 V architectures makes magnetic optimization mission-critical for AI-era infrastructure.

Optimizing Data Center Energy

From legacy server rooms to cloud platforms to AI supercomputing, the world relies on data centers more than ever. This growing demand drives higher energy use each year. In fact, McKinsey projects they will draw 11.7% of all US power consumption by 2030 - up tenfold from 2020. A single NVIDIA GB200 NVL72 rack now draws roughly 120 kW - ten times a conventional CPU rack pushing total facility loads from tens to hundreds of megawatts.

Innovation continues to reshape how organizations store and process data, but most conversations around energy efficiency still focus on cooling systems and renewable power, or on the headline metric of Power Usage Effectiveness (PUE). Less attention is paid to the components that control how electricity moves through the rack, yet these decide how much of the incoming kilowatt actually reaches the silicon.

Magnetics waste a disproportionate amount of energy as heat if not tailored to the task. Across the complex stages of power conversion - from AC/DC front-end rectification through intermediate bus conversion (IBC) down to point-of-load (PoL) regulators - magnetic components can be responsible for up to 80% of conversion losses in poorly optimized designs, and still 25-50% of the residual losses in state-of-the-art 96-98% efficient supplies. Nicola Rosano, the Data Center Application Specialist at Standex Edge, explains: "Designing a magnetic component well can be the difference between 0.5 - 1% loss. If you scale that up, you could save an entire data center hundreds of megawatts of power dissipation at any given moment."

A well-designed transformer boosts efficiency by providing the lowest possible loss while utilizing the smallest possible space. Rosano continues: "Between 0.25% and 0.75% of a circuit's overall power loss can be determined by the design choices of the primary isolation magnetic." Magnetics also shape downstream behavior such as electromagnetic interference (EMI). "If you can't pass EMI-EMC testing, you can't put your product in the market," Rosano adds. It's time the industry leaned on custom magnetics to unlock hidden efficiencies.

The untapped efficiency gap

When space is not limited, there is little efficiency difference between planar and wire-wound magnetics. In real data centers, however, height and volume constraints of typical 1U/2U server slots and the 19" rack pitch shape every design decision. These restrictions drive demand for higher power density solutions. "You're never going to go to [an electronics supplier] and find an off-the-shelf transformer that can do the job," Rosano says, which makes custom design essential. Planar transformers, with their PCB-integrated windings, deliver the low-profile geometry, tight windings, repeatability and low leakage inductance that high-frequency resonant topologies (LLC, PSFB, CLLC) demand - properties that wire-wound parts struggle to match below 10 mm of total height.

Smaller magnetics do more than save space. "The closer components sit together, the more efficient your power supply can be. The more efficient they are, the less EMI problems," Rosano adds. Even in converters already operating at 96-98% efficiency, magnetics account for 25-50% of the remaining losses. Looking at a 3kW PSU, saving just 0.1% efficiency equals 3 W. Multiplied across the ~200,000 PSUs of a single hyperscale data center, that is roughly 0.6 MW saved continuously; across the global server fleet (tens of millions of units), the figure reaches gigawatt range. When you multiply this by hours of operation and cost of energy, the savings can translate into tens of millions of dollars per year and a measurable improvement in facility PUE.

Managing parasitic concerns within magnetic components also improves surrounding circuitry. Lower heat, noise, and ringing allow faster switching and less stress on adjacent field effect transistors (whether silicon MOSFETs, GaN HEMTs or SiC MOSFETs). This can support longer lifespans and more stable operation.

In advanced power electronics, designers should treat parasitics as part of the circuit, rather than as defects to suppress. Rosano explains: "To prevent the elements needed for resonance in complicated topologies [such as LLC and phase-shifted full-bridge converters] from inducing losses into the windings of the transformer, steps must be taken to integrate parasitic elements while minimizing thermal loss." By controlling these parameters within tight tolerances (leakage inductance to +/- 5%, interwinding capacitance below a few hundres pF), systems run cooler, cleaner, and more efficiently.

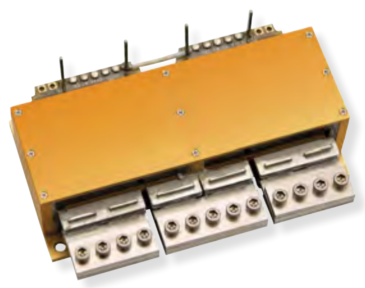

Driving Data Center evolution

As data center power-delivery architectures evolve, OEMs must adapt quickly. The industry is rapidly moving from legacy 12 V intermediate buses to 48 V (Open Compute Project, OCP) and, for AI clusters, to 800 V DC distribution under the OCP ORv3 / Mt. Diablo specifications. Each architectural step changes the turns ratio, isolation rating and core volume required from every transformer in the rack. At Standex Edge, we know early innovation between our component specialists and our clients' system designers shortens development cycles and helps avoid late-stage redesigns. Effective design also means considering downstream effects on the entire system. Material availability, sourcing constraints, and lifetime operation costs all shape real-world performance. Rather than evaluating components only at unit price, we work with OEMs to assess how design decisions affect total cost of ownership over the lifespan of a data center. "Working with so many industries gives us a valuable perspective," Rosano explains. "Instead of seeing three of four power supply designs a year, we see thousands. We see what's state-of-the-art, what everybody's doing. We're the interconnecting point to help connect the dots between typically disassociated people."

One of the biggest shifts is the rise in switching frequencies at the AC/DC front-end and DC/DC isolation stages, towards the megahertz range (1-7 MHz is now common for LLC converters above 3 kW). Point-of-load regulators feeding GPUs and ASICs remain in the 0.5-2 MHz window but are now delivering 1000 A+ at sub-volt rails, a regime where every nanohenry of inductor parasitic counts. As this continues, material choice and construction methods become essential. Low-loss MnZn ferrites (3F36, 3F46, ML91S), nanocrystalline and powder alloys, alternative core geometries (ER, PQ, planar matrix), and wide-bandgap semiconductors - both gallium nitride (GaN) for sub-650 V stages and silicon carbide (SiC) for 1.2 kV-3.3 kV front-ends will mean devices can work at higher voltages, temperatures, and frequencies with lower power loss.

Coordinating suppliers worldwide for specialized applications brings its own challenges. Aligning design, sourcing, and production is critical to meeting time-to-market expectations in fast-moving sectors such as cloud infrastructure and generative-AI deployment, where rack designs are refreshed on 12-18 month cycles. Standex Edge's role increasingly sits at this intersection.

Every OEM should be considering how to balance design priorities such as efficiency, reliability, and cost. No single parameter can be optimized in isolation. Rosano explains: "[Standex Edge's] goal is to find the sweet spot, from not only our components, but to guide our partners to make those best choices." This ensures the best architectural choices, with every micro-component pushing the boundaries of what's possible in data center performance.

The impact of AI

AI adoption is putting unprecedented demand on data centers. According to the International Energy Agency, it's driving up electricity use in accelerated servers by 30% per year. This growth increases pressure on power-conversion efficiency and on

architecture tuned for AI processing. A modern GPU training cluster can draw more power in a single rack than an entire row of CPU servers a decade ago, which is why incremental percentage gains in magnetic efficiency translate directly into deployable compute density.

At the same time, AI will likely play an increasing role in transformer design. As in other markets, there are opportunities for AI to streamline the iterative loops of magnetic design winding configuration, Dowell-based AC resistance optimization, core loss prediction with improved Generalized Steinmetz Equation (iGSE) models, and finite-element-based leakage extraction. "Engineering expertise won't be replaced, but we want AI to help find our customers a faster, better, and higher fidelity solution, which is ultimately our end goal."

These tools can shorten design cycles. Rosano explains: "Standex Edge is actively deploying R&D, commercial and engineering resources to figure out how to grow alongside this market and develop differentiated products." However, the experience and judgment of our specialists remain essential to ensure designs meet safety, reliability, and production requirements. This approach will help Standex Edge stay ahead of changing data center demands.

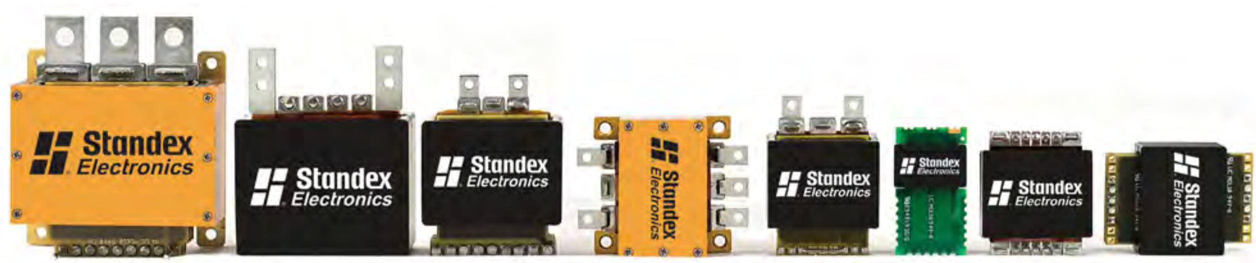

Key takeaways

Magnetics are the silent bottlneck:
In a 96-98% efficient PSU, 25-50% of residual losses still come from the transformers and inductors.

Off-the-shelf does not scale:
1U/2U height limits, 48 V and 800 V bus migrations and MHz switching make every winning design a custom planar.

Parasitics are circuit elements, not defects:
Tight tolerances on leakage inductance and interwinding capacitance let LLC, PSFB and CLLC topologies hit their resonance points cleanly.

0.1% matters:
Three watts saved per PSU x 200,000 PSUs per hyperscale x 8,760h x $0.08/kWh = ~$420,000 per year, per data center.

AI is both the load and the design tool:
AI servers drive demand; AI-assisted magnetic design helps meet it.

A Standex **Electronics** Business

The pathway to usable electrical power begins with generation, continues through high-voltage transmission to urban and industrial centers, and extends to the distribution of lower-voltage electricity to businesses and homes, where it's ultimately consumed in countless applications.

At the **Edge** of this energy cycle, Standex Electronics engineers open up a virtually infinite range of possibilities to power the world's most demanding electrical devices. From concept to production, we deliver custom magnetic solutions that convert, condition, and control energy with precision and reliability.

Our expertise spans transformers, inductors, current sensors, and other advanced magnetic components used across industries, from automotive and industrial to medical and aerospace, including the high-frequency planar transformers and matrix-array architectures now powering hyperscale and AI data centers worldwide. Combining decades of engineering innovation with global manufacturing capabilities, Standex Electronics stands at the **Edge** of innovation, where energy meets possibility.

Welcome to **Standex Edge**.

standexedge.com

EDGE-COM-WP260501