HOW2POWER TODAY
Your Power Design Newsletter

Power Products

ISSUE: June 2026

Vertical Power Delivery Platform Tackles Power And Thermal Limits

Lotus Microsystems' vStrata is described as the the first low-profile platform designed to solve electrical, thermal and mechanical constraints in power delivery for AI processors as a single, co-engineered system. vStrata unifies electrical, thermal and mechanical design into a single architecture. Built on a silicon-based substrate with superior thermal conductivity, vStrata manages heat directly at the point of load—eliminating hotspots and reducing operating temperatures by up to 25°C in optimized configurations, according to the vendor (see the figure).

As modern AI accelerators draw increasingly massive current loads, legacy architectures are struggling to keep pace. The industry is shifting toward vertical power delivery (VPD) to solve last-inch distribution losses, but existing solutions treat power and thermal challenges as separate problems, says the vendor. vStrata solutions address both simultaneously using proprietary silicon power interposer technology (PIT), bringing power directly beneath the processor while managing thermal load at the same point.

"We have reached a point where AI compute performance is constrained by physical architecture," said Hans Hasselby-Andersen, CEO of Lotus Microsystems. "vStrata was developed to address the reality that power delivery and thermal management have become inseparable system constraints. By integrating both into a single co-engineered platform, we help data centers increase compute density while reducing the power and cooling overhead that drives AI infrastructure cost."

Engineered for kiloampere-class AI power demands, the LSC0580 achieves up to ~96% point-of-load efficiency. By simultaneously reducing last-inch power losses and thermal bottlenecks, the platform lowers power and cooling overhead while enabling higher sustained xPU performance. In addition, proprietary silicon PIT and advanced packaging technology enable ultra-thin architectures, with a roadmap below 1 mm. By minimizing electrical distance between power and processor, the platform improves dynamic performance and supports load transients exceeding 10 A/ns without external capacitors.

Designed for integration with established Tier-1 reference designs and existing power management controllers, vStrata reduces adoption friction while enabling next-generation performance, says the vendor. The LSC0580 has successfully taped out for leading xPU and AI infrastructure partners.

"The industry has been trying to scale AI compute using an architecture that was never designed for this level of intensity," said Yasser Nour, co-founder and CTO of Lotus Microsystems. "The real bottleneck isn't just delivering power—it's how the system responds to rapid, unpredictable load changes. By collapsing the distance between the power supply and the processor, vStrata is designed to deliver current where and when it's needed, without compromising thermal behavior or transient stability."

Engineering samples of vStrata LSC0580—the first platform module—will ship in Q3 2026. The company is currently working with Tier-1 hyperscalers and expanding engagement with additional partners globally through its Early Access Program, with the first engineering samples to be delivered in Q3 2026. To learn more, see the vStrata page or contact the company.

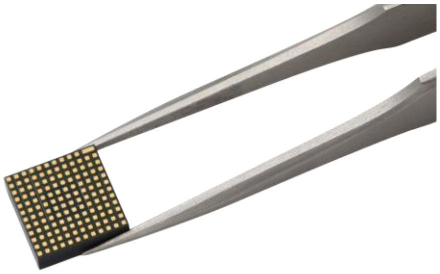

Figure. vStrata is a vertical power delivery platform for AI accelerators that co-designs power delivery and thermal management in a single module. Described as a first-of-its-kind, the platform reduces power conversion losses by more than 50% and lowers AI processor operating temperatures by up to 25°C, according to the vendor.