ISSUE: April 2025

APEC 2025: Industry Crafts Vertical Power Delivery Solutions For AI Processors

by David G. Morrison, Editor, How2Power.com

Continuing a trend that that began at last year's conferences, if not earlier, AI was once again the major topic of discussion at this year's APEC in Atlanta. Suppliers of power components and power modules, spoke on power requirements in AI data centers, and showcased their solutions with emphasis on power components for vertical power delivery to GPU and CPU processors.

Meanwhile, wide-bandgap SiC and GaN devices were also featured in many APEC exhibits, as they have been for many years now. However, much of the news concerning these devices was packaging related. In terms of actual semiconductor developments, one major supplier previewed its introduction of a SiC JFET for solid-state circuit breakers (SSCBs)—giving the industry another source for these components. And there was other product news relating to the SSCB applications.

While AI remained the dominant application under discussion in the exhibition, robotics was also a popular theme. In what's become a routine occurrence at recent electronics conferences, one vendor was seen walking a robot through the expo—this time a humanoid robot. Like the robot dog at the vendor's booth, this robot demo was meant to highlight the need for GaN power devices in driving the many motorized joints in robot applications.

This article offers examples of some of the developments described above along with other product news from APEC 2025. However, in light of recent industry news, namely the layoffs that have been occurring in recent months at so many power semiconductor companies, it's worth noting the business context in which these product developments were discussed. Despite the enthusiasm around AI datacenters, and to a lesser extent robotics and a few other areas, the electronics industry is in a slump.

As one semiconductor industry veteran observed, "The unspoken theme of APEC was that the industry is down and has not recovered from Covid stockpiling of inventory. Worse still, nobody knows when the demand for semiconductors (outside of AI chips and systems) will recover."

Powering AI

One week before APEC, Infineon announced its latest generation of high-density power modules for AI and high-performance computing, the OptiMOS TDM2454xx quad-phase power modules. These modules were on display for the first time at Infineon's booth and discussed at the company's annual APEC press dinner.

The TDM2454xx modules enable true vertical power delivery (VPD) and are said to offer an industry-best current density of 2 A/mm². The modules follow the OptiMOS TDM2254xD and the OptiMOS TDM2354xD dual-phase power modules introduced by Infineon last year. See Figs. 1 and 2.

The OptiMOS TDM2454xx modules are described as a fusion of Infineon's robust OptiMOS 6 trench technology, chip-embedded package for superior electrical and thermal efficiencies, and innovative low-profile magnetic design that continue to push the envelope for performance and quality of VPD systems. Additionally, the OptiMOS TDM2454xx has a footprint that is designed to enable module tiling and improving current flow that enhance electrical, thermal and mechanical performance. According to Infineon, the TDM2454xx are described as the first modules optimized for VPD with a footprint designed for optimal tiling arrays of modules.

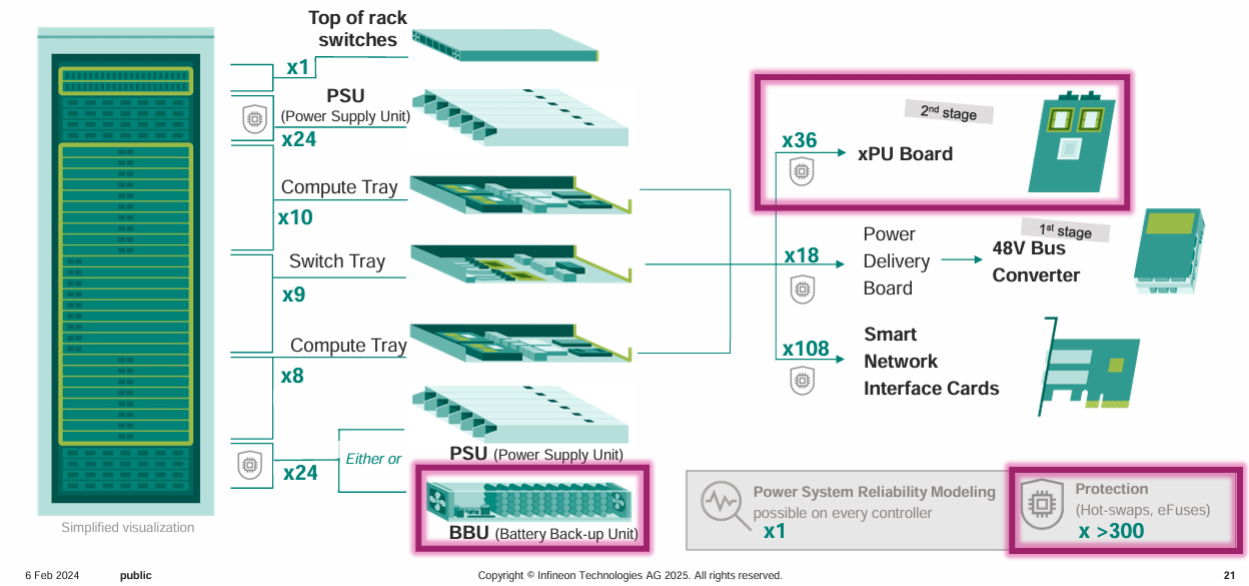

Fig. 1. This slide from Infineon's press dinner presentation at APEC 2025 in Atlanta, highlights the multiple opportunities for power components in AI server racks. Infineon and other companies have been discussing the power supply, BBU and 48-V to 12-V (or lower) bus converter applications extensively at Electronica, and other recent conferences. However, at APEC 2025, vendors seemed particularly focused on components and solutions for the GPU/CPU power stage, which converts 12 V to the ~1 V required to power the processor. Most of the components and solutions shown or demo'd were providing vertical power delivery, which places the power stage components on the bottom side of the board under the GPU/CPU to minimize the interconnection paths between the VR components and the processor.

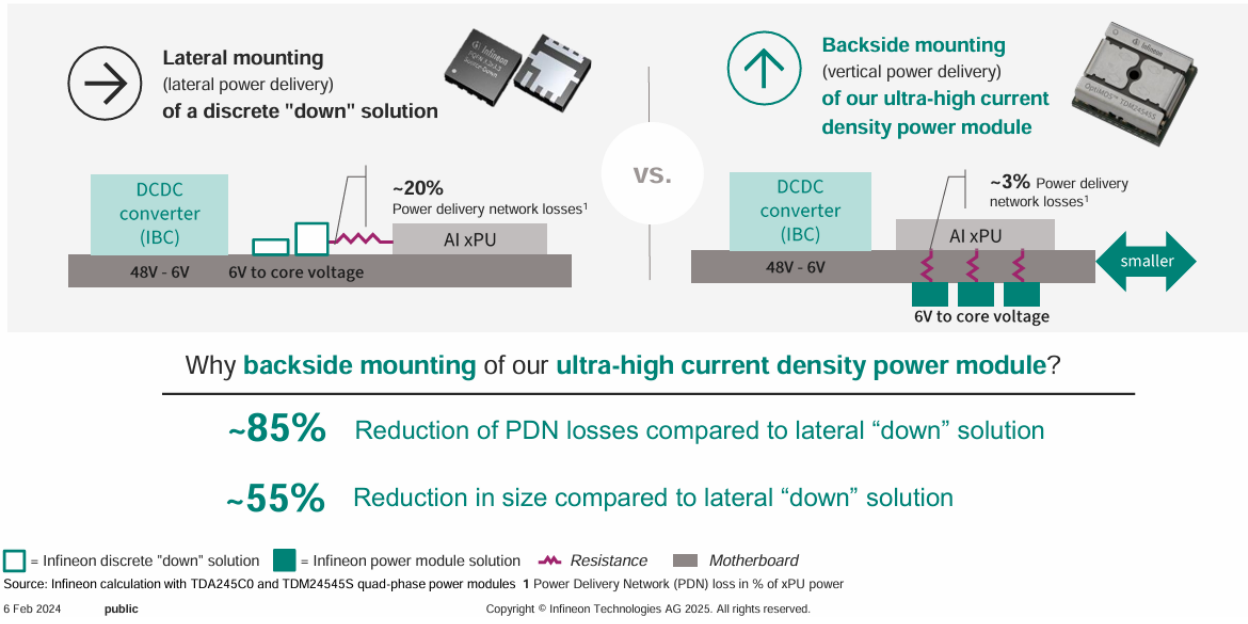

Fig. 2. Infineon's comparison of vertical power delivery with lateral power delivery. In traditional horizontal power delivery systems, power needs to travel across the surface of the semiconductor wafer, which can result in higher resistance and significant power loss. Vertical power delivery minimizes the distance that power needs to travel, thereby reducing resistive losses enabling increased system performance. However, there are tradeoffs with vertical power delivery such as potential requirements for heatsinking on both sides of the board.

The OptiMOS TDM2454xx modules support up to 280 A across four phases with an integrated embedded capacitor layer within a small 10-mm x 9-mm form factor. Besides the embedding of caps, the 12 silicon MOSFET die required to support the four phases at the specified ratings are also embedded within the module's

substrate. As with the previously gen 1 and gen 2 modules, the TDM2454xx incorporate proprietary magnetics—in other words, custom inductors with custom cores (see Fig. 3).

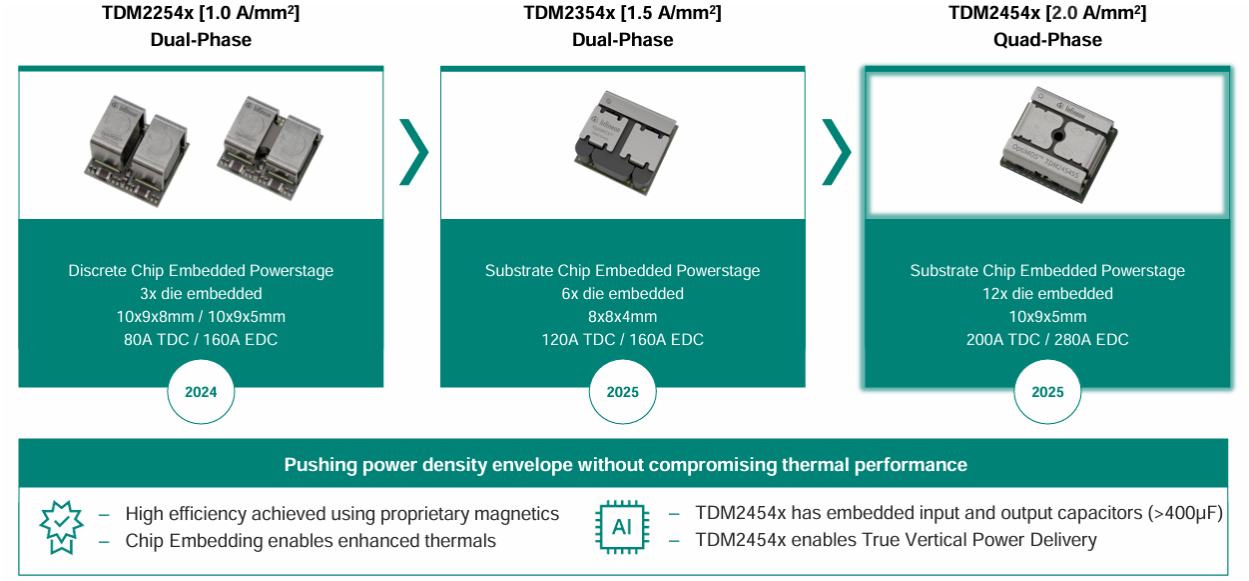

Fig. 3. In a relatively short window of time, Infineon has advanced its power stage technology from a current density of 1 A/mm² in its first-generation power modules to 2 A/mm² in its third-gen power modules.

Combined with Infineon's XDP controllers, the power stage modules are said to offer a robust power solution with improved system power density. For more information, see the OptiMOS quad-phase power modules page.

Interestingly, Texas Instruments also showed power modules for GPUs in its booth. These were two-phase power modules with four MOSFETs and inductors in a 9-mm x 10-mm x 5-mm format. These are the same module dimensions as Infineon's TDM2454x quad-phase power stage. Although TI's modules are not yet released, they illustrate the growing interest in vertical power delivery.

Another company showing power components for vertical power delivery was Empower Semiconductor. In their exhibit, they featured Empower's Crescendo platform, which is based on the company's proprietary FinFast technology and integrates all power components into a single ultra-thin device, enabling mounting of these devices beneath the processor, while eliminating the requirement for a large bank of decoupling capacitors.

According to the company, this breakthrough shrinks the AI power supply by 5x and delivers on-demand kilowatt power with unmatched speed and accuracy while reducing power distribution losses by 10% to 20% (or more) versus a traditional lateral power delivery scheme (Fig. 4).

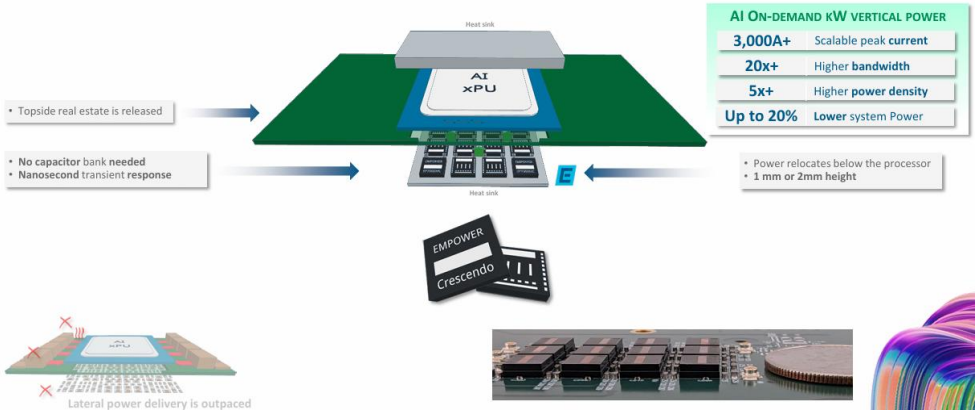

Fig. 4. Within a custom power, LGA-style package standing 1 or 2 mm high, Crescendo devices integrate very high speed voltage regulators employing FinFET cells as power devices together with advanced magnetics and silicon capacitors.

"Crescendo tackles one of the biggest hurdles in AI-driven data centers—how to deliver the increasing power demands of new processors with the efficiency, responsiveness, and compact footprint required to support their unprecedented performance," said Empower's CEO and founder, Tim Phillips. "We have been partnering with industry leaders on multiple platforms since announcing its availability for design last October and are gearing up for production by the end of the year. We're coming to APEC with a live demo and presentation that highlight the breakthroughs we have achieved."

Mukund Krishna, senior manager, Product Marketing at Empower Semiconductor explains that the high switching speed and the resulting 20X bandwidth of the company's integrated voltage regulators, which use Finfet power switches, enable the in-package capacitors of the processor being powered to provide the necessary decoupling. Consequently, the Crescendo solution does not require all the ceramic caps on the processor board that would be needed with other power solutions for vertical power delivery.

A single IVR which integrates inductors and silicon e-caps from Empower can deliver 65 A and up to 50 of them can be combined to deliver ~3 kA as depicted above in Fig. 4. As a demo at APEC, the company's showed a 4 x 3 array of these devices capable of 800-A output, which would be on the low end of requirements for upcoming AI processors.

While this solution was announced in October at the OCP Summit, at APEC the company showed the full solution live under load step conditions. The company is targeting end of year for release of the IVR and its supervisory chip, which will then complete the Crescendo solution.

In their exhibit, Ferric, a start-up company based in New York city, showcased its very high density integrated voltage regulator (IVR) solutions and a demo of its AI power solution, featuring chip-scale power converters that are described as 25 times smaller, >100 times faster, and >30% more efficient than alternatives.

The company's first part is the Fe1736, a 16-phase buck converter with monolithically integrated power switches and co-packaged inductors in a µLGA that measures just 5.64 x 3.54 x 0.6 mm. Operating from a 1.8-V input, this point-of-load converter generates a 0.25-V to 1.5-V output at up to 56 A.

According to Ferric's CEO, Noah Sturcken, the buck converter switches with a peak efficiency of >91% at an astounding 30 MHz per phase, which leads to an even higher effective switching frequency given the 16 phases. Some designers may opt to run the converter at even higher switching frequencies (up to 100 MHz) to improve transient response. The demo board displayed at APEC featured an array of these devices configured to produce almost 1800 A for powering an AI type processor (Fig. 5).

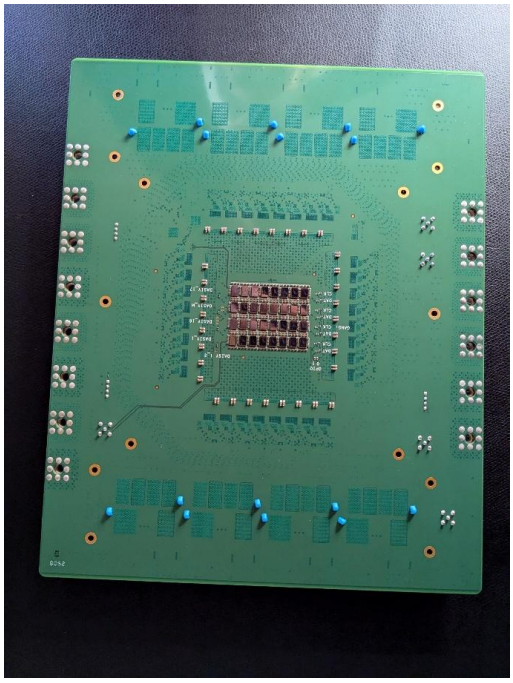

Fig.5. An array of FE1736 IVRs may be mounted to the backside of a PCB to power a GPU or other processor using vertical power delivery. Beyond that, they are even small enough to be mounted on the die side of the processor substrate or the land side of the package.

These capabilities were made possible by the company's development of the proprietary thin film materials it uses to build inductors on wafer. Indeed, the company had its origins in teaching TSMC how to build inductors on wafer, and in CEO's doctoral research on ferromagnetic materials, says Sturcken. He notes that the development of these core materials revolved around the tradeoff of flux density and loss.

The Fe1736 is currently sampling with production quantities expected to become available in Q1 2026. A quad-output version of this regulator, the Fe1728, is also sampling.

Addressing Other AI Data Center Power Requirements

As the graphic in Fig. 1 above illustrates, data centers require bulk power supplies to convert ac power to 48 V, bus converters that stepdown 48 V to 12 V (or lower voltage) and battery backup units (BBUs). In addition to discussing the new power modules for VPD at APEC, Infineon also shared information on its 12-kW BBU solution.

Infineon's proprietary topology for BBUs is based on partial power conversion, which is said to enable smaller magnetics and a more compact design (Fig. 6). According to the vendor, it outperforms existing solutions in terms of efficiency, power density and cost-effectiveness.

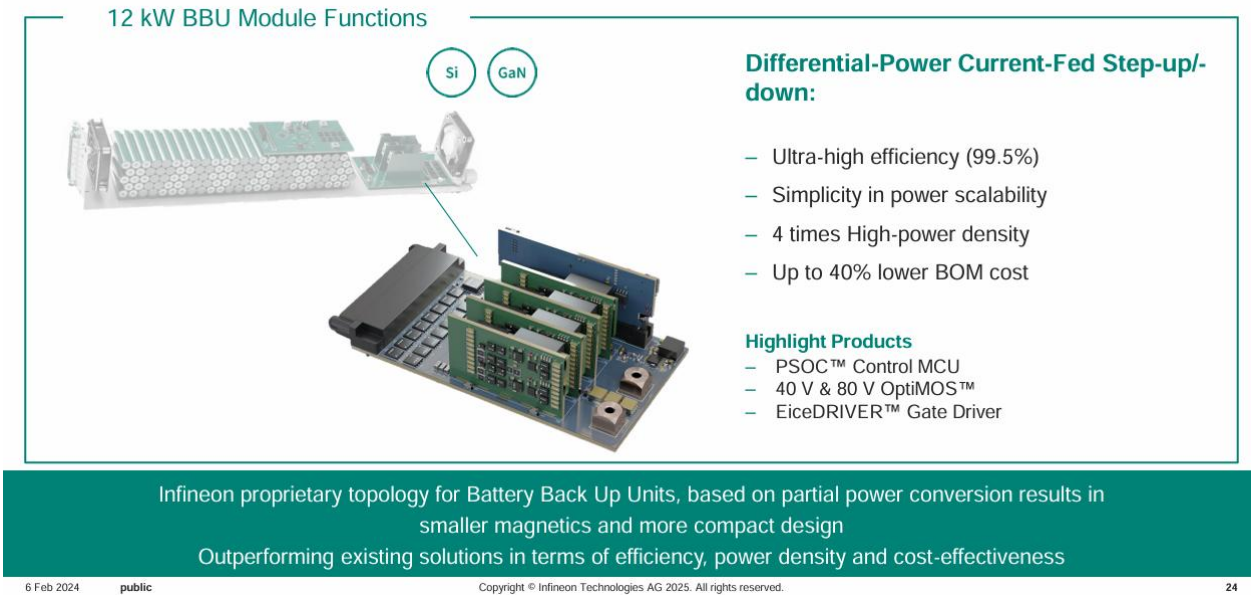

Fig. 6. Infineon's latest 12-kW battery backup unit (BBU) solution addresses increasing power demands and strict space constraints. Efficiency of the BBU is becoming especially important in AI applications where it is used not only when the power source fails, but also to assist the server power supply by providing peak shaving during periods of high load demand. With that efficiency requirement in mind, Infineon applies its partial power conversion topology to the BBU, whereby the BBU's power converter adds or subtracts a small voltage to the battery voltage to obtain a regulated 48-V output. In this way, the BBU's losses are reduced and thus its efficiency is higher.

At APEC, Infineon also unveiled a new member of its XDP XDP7xx family of wide input voltage range digital hot swap controllers—see the "48-V Digital Hot-Swap Controller Supports High Power For AI Servers" in this issue of the newsletter.

Also see news of a related device from Texas instruments, the TPS1685 hot-swap eFuse, discussed in "Hot-Swap eFuse And GaN Power Stages Support Data Center Applications". That article contains news of the company's new 650-V GaN FET power stages which target server power supplies among other applications. These devices were introduced at APEC 2025 where TI showcased additional power components and reference designs for AI data centers. Among these were a a 3.6-kW CRPS power supply reference design, which employed 50-V GaN in the TOLL package and 100-V GaN for sync rectification.

Additionally, TI showed an open-rack PSU reference design for data centers. Measuring 32-mm x 68.5 mm x 700-mm, the 8-kW PSU achieves high power density and a high efficiency of 97.5% (targeting Titanium level)

The design features an interleaved three-phase TCM TP PFC and three-phase delta-delta LLC, leveraging GaN devices in both stages. Operating from universal ac input it generates 48-V output.

Another data center-related exhibit from TI was an 8.1-kW OCP-V3 battery backup unit (400-V battery to 48 V). This unit contains a 3.6-kW bidirectional dc-dc converter based on a DAB topology. Bosheng Sun, senior systems engineer at Texas Instruments, observed that this BBU operates totally under digital control.

At the EPC booth, the company's interest in AI data centers was much in evidence as their booth featured a seemingly life-sized photo of a server rack within a data center. As CEO Alex Lidow observed during my visit to the EPC booth, data centers continue to be the biggest application for EPC's GaN products with fully half of the company's business in servers.

Opening a drawer in its server-like wall display, Lidow pulled out some examples of server power applications for its GaN power devices (Fig. 7). One was a 48-V to 12-V LLC-based bus converter, which Lidow described as the main data center application for the company's products. As implemented in their EPC9159 reference design, the LLC consists of a primary side full bridge, a fixed ratio 3:1 planar transformer, and a center tab synchronous rectifier for the secondary side. The primary comprises four 80-V EPC2619 GaN FETs, while the secondary has six 40-V EPC2067 GaN FETs. The converter is realized with a planar transformer design with custom magnetic cores.

The LLC power stage components are depicted schematically in Fig. 7b and along with a modular implementation of the magnetics and secondary-side FETs. This LLC design has a 3:1 stepdown ratio in through mode but can be configured for a 4:1 stepdown in partial power mode. As part of the full reference design implemented in the EPC9159KIT demo board, this LLC converter delivers 83 A at 12 V with a full load efficiency >95.5%. When assembled as a complete LLC converter module, the design is said to enable >5 kW/in.[3].

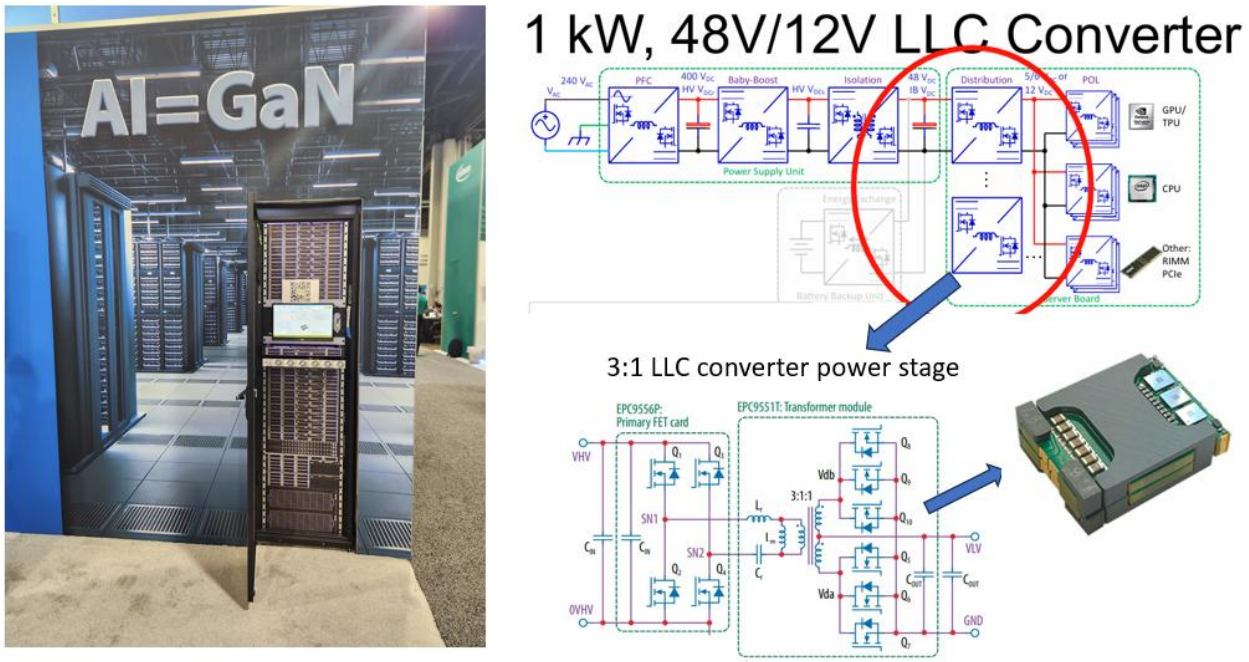

(a) *(b)*

Fig. 7. EPC's life-like data center display (a) at APEC 2025 showcased the presence of EPC GaN components in server power applications such as the 48-V to 12-V bus converters diagrammed in (b). Their EPC9159 reference design for a 3:1/4:1 LLC converter employs four EPC2619 GaN FETs on the primary side, and the six EPC2067 GaN FETs on the secondary. As a demonstrator, EPC has implemented the transformer module portion of the power stage as a 23.8 x 17.5 x 7.2-mm module.

During my booth visit, Lidow also showed a 5-kW OCP-style, ac-dc server power supply that uses EPC GaN in two novel stages: one is a multi-level totem-pole PFC stage that stacks six GaN FETs and the other is an isolated dc-dc stage that stacks four, GaN-based LLC modules (Fig. 8). This GaN-based server power supply pushes power density above 100 W/in.[3].

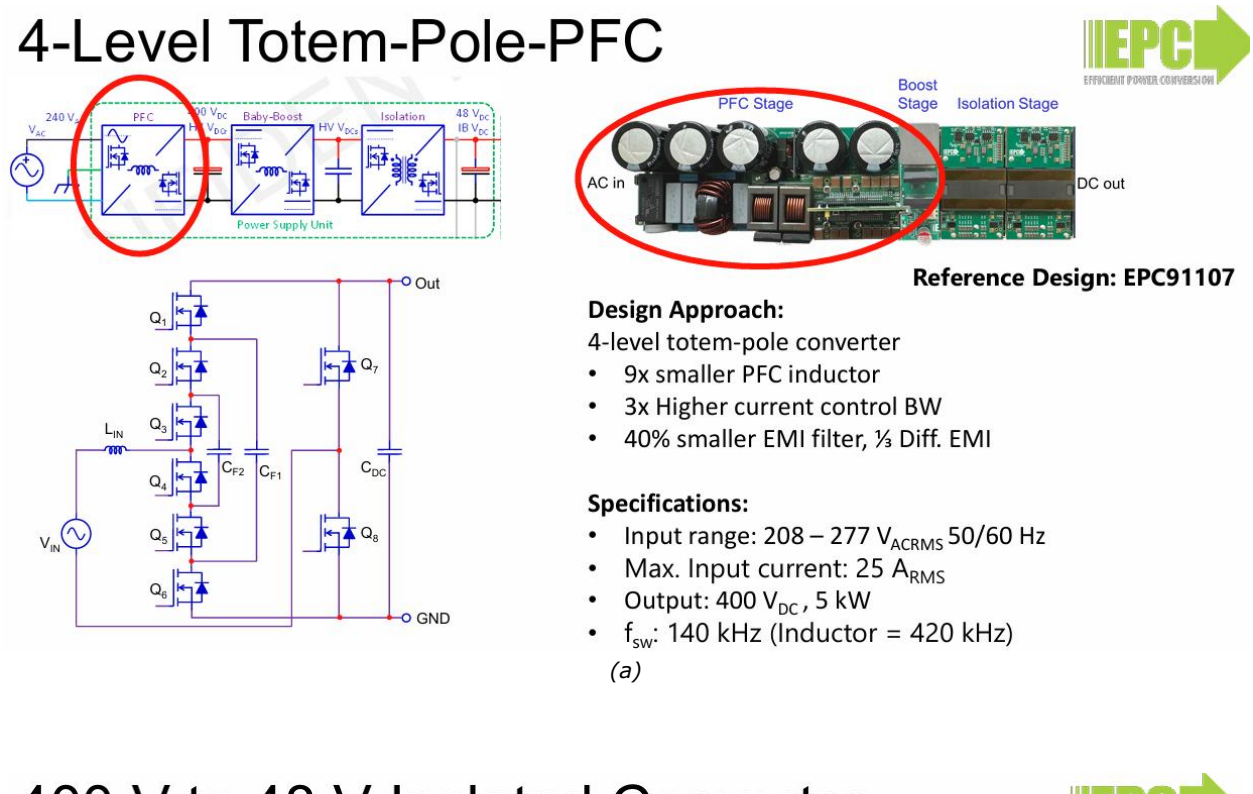

Fig. 8. EPC's implementation of ac-dc server power supply exploits GaN performance in novel PFC and isolation stages. A GaN-based TP-PFC stage stacks six 200-V EPC2304 GaN FETs in a multilevel topology that shrinks magnetics significantly. Meanwhile, the isolated converter stacks four GaN-based LLC stage modules to step down the 400-V bus to 48 V. The complete 5-kW power supply pictured in the upper right of (a) and (b) measures 275 x 70 x 40 mm, which (pardon my English) corresponds to 107 W/in.[3]. The baby-boost stage is a simple boost converter that provides voltage regulation for the isolation stage, or extra hold-up time, as needed.

Powering Robots

Calling attention to its growing participation in the robotics market, EPC has augmented its booth staff at recent events with robot companions. At previous shows I met their robot dog Chip. This year, I also met a humanoid

robot named Greg, a slim skeletal looking creature that could be seen marching though the exhibition hall, sometimes in a robot themed t-shirt (Fig. 9a).

Meanwhile, back at the EPC booth, their displays revealed more about the role of GaN in humanoid robots. Lidow held up a motor drive reference design implemented on a small circular. PCB that combined three EPC23104 GaN power stages with two current sensors and an MCU within a 32-mm diameter circle (Fig. 9b). The GaN FETs enable the small size of the inverter, which then fits within the many joints of the robots.

So many of the EPC23104 power stages are being used in these applications, "we can't keep them on the shelf," said Lidow who also noted that GaN accounts for about $300 worth of content in a humanoid robot.

(a) (b)

Fig. 9. Greg the humanoid robot (a) served as EPC's roving ambassador for GaN in the growing robotics market. A motor drive reference design (b) illustrates why GaN is gaining traction in these applications. Using GaN power stages together with current sensors and MCU for control, a complete inverter can be implemented on a circular board measuring 32-mm across. That's all within the inner circle shown on the PCB pictured. The full board has a diameter of 55 mm, which was the size of the original inverter design using silicon MOSFETs. This particular GaN inverter drives the humanoid's wrist motor, but there are numerous joints throughout the robot that may require similarly small motor drives.

Packaging Developments

At its booth in the exhibit hall, Infineon displayed demo boards showing parts in some of the company's new packages. For example, Infineon is introducing members of its 1200-V CoolSiC MOSFET product line in Q-DPAK packages. This package style is notable for its use of top-side cooling. The package is said to feature an innovative molding compound that reduces the creepage requirement from 6 mm to 4.8 mm. Additionally, the company's .XT diffusion soldering of the chip to a copper rather than aluminum leadframe improves thermal performance of the device.

The company will offer different device configurations (single and multi-MOSFET devices) and die sizes in the Q-DPAK as identified in the chart (Fig. 10). The 1200-V devices join the 650-V devices already offered in this package, and the new 1200-V parts will cover a range of $R_{DS(ON)}$ values from 4 to 78 mΩ. So as Michael Williams, director of product marketing, Industrial and Infrastructure at Infineon Americas, observed, CoolSiC in Q-DPAKs are driving scalability by both voltage rating and on-resistance. Said Williams, the QDPAK can hold 1 or 2 die, which enables it to achieve best in class $R_{DS(ON)}$ in a package of just 2.3-mm height. Typical applications for these 1200-V devices include dc-dc converters and on-board chargers in EVs.

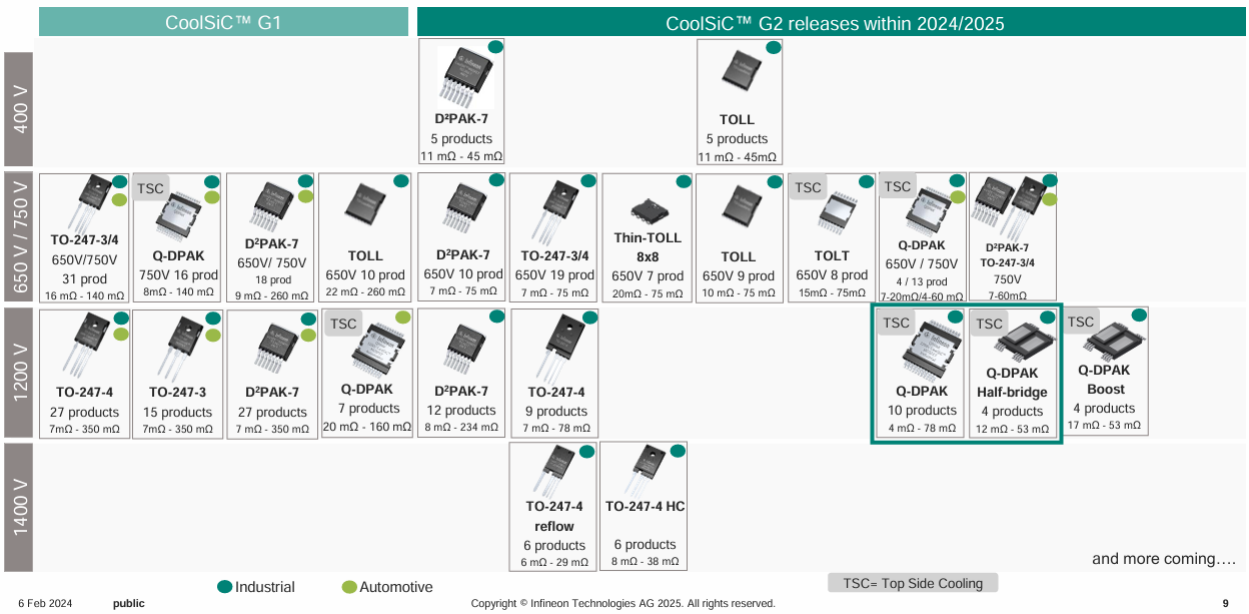

Fig. 10. Infineon's near-term roll-out of CoolSiC MOSFET G2 discretes, industrial and automotive grade: 400 V, 650 V, 750V, 1200 V and 1400 V. The latest parts include 1200-V devices in Q-DPAKs.

Infineon also exhibited its CiPOS Maxi IPMs (intelligent power modules (IPMs)). The company offers both IGBTs and CoolSiC MOSFETs in these packages including 1200-V SiC devices. These devices currently target motor drive applications, but in the future will be offered for dc-dc converters and OBCs in EVs and for renewable energy applications, said Williams, who added that devices in several topologies will be introduced this summer.

In addition, at the press dinner, Infineon discussed its new Easy 2 series packaging for power modules (Fig. 11). This package is constructed using a new plastic material and new silicon gel that permits 175°C steady-state operation and 200°C overload. It also features new press-fit pins with 2X the current capacity of previous pins. It's currently offered with an Al_2O_3 (alumina) substrate but will be offered with AlN (aluminum nitride) substrate in the future to accommodate different application requirements.

The company will offer its 1200-V G2 CoolSiC devices in this package. EV charging, in particular is said to need the high performance of the Easy 2 series power modules.

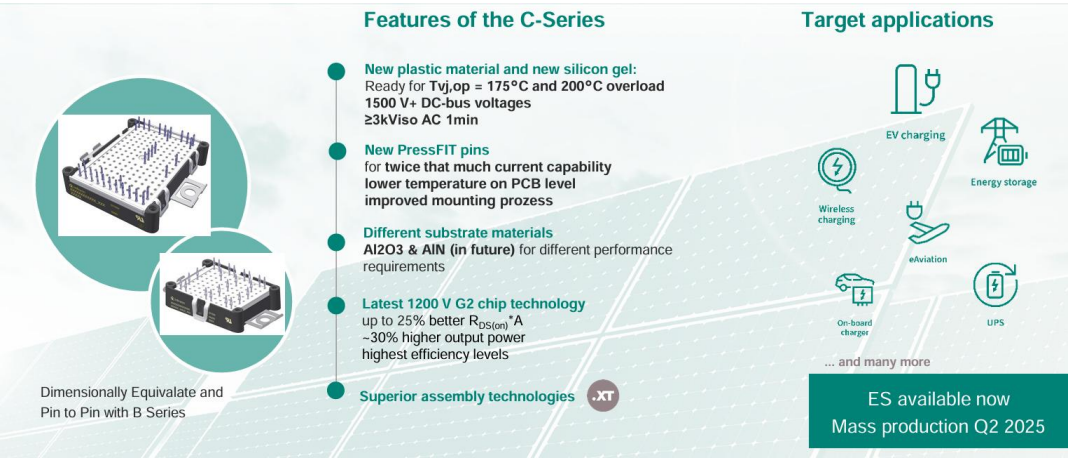

Fig. 11. Infineon's Easy 2 C-series will offer the company's latest G2 CoolSiC semiconductors in an advanced version of the industry-standard Easy package, which leverages new housing materials and metallization to deliver higher power performance.

At the Nexperia booth, Art Gonsky, senior principal product engineer, discussed the company's introduction of 1200-V SiC MOSFETs in the top-side cooled X.PAK. This surface-mount package, with its compact form factor of 14 mm x 18.5 mm, combines the assembly benefits of SMD with the cooling efficiency of through-hole

technology, ensuring optimal heat dissipation. In addition to its thermal and automated assembly benefits, the X.PAK package enables low inductance for surface-mount components.

In the formal announcement for this part, said Katrin Feurle, senior director and head of SiC Discretes & Modules at Nexperia, said, "The introduction of our SiC MOSFETs in X.PAK packaging marks a significant advancement in thermal management and power density for high-power applications. This new top-side cooled product option builds on our successful launches of discrete SiC MOSFETs in TO-247 and SMD D2PAK-7 packages. It underscores Nexperia's commitment to providing our customers with the most advanced and flexible portfolio to meet their evolving design needs."

Gonsky added, "Nexperia's SiC MOSFETs in X.PAK package offer a compelling power-to-size ratio, driven by the inherent advantages of SiC technology and the X.PAK package design. These packages can handle high power levels within a compact footprint, leading to increased power density and reduced system size. There's nothing close to them in power."

These devices are targeting industrial applications such as battery energy storage systems (BESSs), photovoltaic inverters, motor drives, and UPSs. Additionally, they are well-suited for electric vehicle charging infrastructure, including charge piles. The initial portfolio includes products with $R_{DS(ON)}$ values of 30, 40, 60 mΩ (NSF030120T2A0, NSF040120T2A1, NSF060120T2A0) and a part with 17 mΩ has been recenty released. An automotive-qualified SiC MOSFETs portfolio in X.PAK packaging will follow later in 2025, as well as additional $R_{DS(ON)}$ classes like 80 mΩ. For more information, see the SiC MOSFETs page.

Solid-State Circuit Breakers

At the press dinner, Infineon's vice president of SiC, Peter Friedrichs, previewed the company's first JFETs, which will include a 1.5-mΩ, 750-V device and a 2.3-mΩ, 1200-V device that will target solid-state circuit breaker applications. Previously, perhaps the only vendor who was marketing and publicizing SiC JFETs was United SiC. (That company was acquired by Qorvo in 2021, but Qorvo recently sold its United SiC product line to onsemi.)

SiC JFETs are considered particularly well suited for SSCBs because, unlike most power MOSFETs, they are depletion-mode devices and therefore normally on. In addition, they offer especially low on-resistance for a given die size—lower than a similarly rated SiC MOSFET with comparable die size. Friedrich noted that the Infineon's JFETs will be especially optimized for SSCBs rather than for power supply switching applications as may have been the case with some previously introduced devices. The company plans to release more information on the new JFETs at the upcoming PCIM.

Other Product Developments Seen At APEC 2025

- Power Integrations' launch of TinySwitch-5. See "Fifth-Generation Switcher Family Delivers 175 W And 92% Efficiency In Flybacks".
- Nexperia has added 12 new devices to its continuously expanding e-mode GaN FET portfolio.
- Texas Instruments displayed a QR flyback design based on the UCG28226 converter IC. It's self biasing, taking its power from the switch node and eliminating the need for an aux winding on the transformer.
- Texas Instruments also showed a wired BMS for energy storage (BESSs). It was comprised of three reference design boards. The TIDA-010279 battery management unit, which performs cell monitoring and cell balancing; the TIDA-010272 rack monitor unit, which performs pack monitoring of current and temperature and can sync the boards for fuel gauging algorithms; and TIDA-010253 battery control unit, which has an MCU that controls the whole BMS. All the devices used in these designs are functional safety compliant. Overall, the interface is a differentiator—the boards are daisy chained in a ring architecture, such that each board can connect either "north or south" in the event one of the boards it connects to breaks. For more info, see the High-voltage BMS for energy storage systems demonstration video.
- pSemi introduced two new regulators: the PE24111 two-stage buck regulator and PE25213 ultra-high efficiency charge pump. It also previewed what it describes as "the world's first 4-level buck converter for battery charging applications." For more on the PE24111 and PE25213, see "20-A Buck Regulator And 10-A Charge Pump Offer High Efficiency And Low Profile" in this issue.
- Danisense launched a new Transducer Electronic Datasheet (TEDS) functionality.

For Further Reading

While the EV market is somewhat down now, GaN power semiconductor suppliers continue to discuss how their technology is poised to overtake silicon carbide in traction inverters and other EV power conversion applications. This is one of the issues that were discussed in a SiC vs GaN rap session at APEC 2025, moderated by Power Electronics News editor-in-chief Maurizio Di Paolo Emilio. Aalyia Shaukat wrote a very nice report on this session. See "The Great Debate at APEC 2025: GaN vs. SiC".