




POWER ELECTRONICS

ISSUE V 2026

 AN ANGEL BUSINESS COMMUNICATIONS PUBLICATION

POWERELECTRONICSMAGAZINE.NET



RETHINKING AI DATA CENTRE COOLING



Thin film capabilities **reducing energy consumption** enabling new technologies

Emerging Technologies like Augmented Reality will connect us with our friends and give us high speed access to vital information just when we need it. Evatec thin film technologies will deliver the essential high performance displays, connectivity, electronics packaging and battery technologies that all reduce energy consumption in a power hungry world.



Evatec thin film technology solutions, helping to shape a sustainable future.

More about “Energy Efficiency”

Learn more about Evatec production solutions for waveguides and light engines in Augmented Reality.

Read more about Evatec know-how in photonics essential for camera and smart sensing technologies.





The systems era has arrived

➤ If this year's ANGELTECH conference demonstrated anything, it is that power electronics has reached an inflection point. The industry's most pressing challenges are no longer centred on individual device performance, but on how technologies are integrated into complete systems capable of delivering greater efficiency, resilience and scalability. Whether the discussion focused on electric vehicles, AI data centres, wide band gap semiconductors or future energy infrastructure, the same conclusion emerged repeatedly: systems engineering has become the new competitive frontier.

This marks an important shift. For decades, progress in power electronics was measured by incremental improvements in semiconductor devices. Today's priorities are far broader. Device innovation remains essential, but its value is increasingly realised only when combined with advances in packaging, thermal management, software control, manufacturing and power architecture. Performance is no longer defined by the transistor alone, but by the optimisation of the entire energy conversion chain.

Nowhere was this more evident than in the automotive sessions. Electrification continues to gather pace, but the conversation has matured beyond the simple debate over battery electric vehicles versus internal combustion engines. Manufacturers are navigating a far more complex landscape, balancing hybrids, battery electric platforms, charging infrastructure, cost pressures and regional policy objectives. Throughout these discussions, power electronics emerged as the technology that underpins every strategic decision. The continued adoption of silicon carbide and gallium nitride is enabling higher-voltage architectures, faster charging and improved efficiency, but equally important is the industry's growing focus on integrating these devices into robust, manufacturable vehicle platforms.

The rapid evolution of AI infrastructure reinforces this systems-level perspective. As data centres move towards rack densities exceeding 100 kW, conventional approaches to power delivery and cooling are rapidly becoming obsolete. The challenge is no longer one of simply supplying more electricity, but of orchestrating power conversion, thermal management, batteries and software into a coherent, intelligent infrastructure. In this respect, the AI revolution is proving to be as much an energy challenge as it is a computing challenge.

Wide band gap technologies featured prominently throughout the conference, yet perhaps their greatest significance lies in the fact that they are no longer viewed as emerging technologies. Silicon carbide and gallium nitride have entered commercial reality across automotive, industrial, renewable energy and digital infrastructure markets. The industry's attention is therefore turning towards manufacturing scale, reliability, packaging innovation and supply chain resilience—the practical considerations that ultimately determine commercial success.

Equally striking was the growing convergence between sectors that have historically operated independently. Automotive manufacturers are increasingly concerned with grid capacity. Data centre operators are becoming sophisticated energy managers. Utilities are embracing advanced semiconductor technologies to improve grid flexibility and renewable integration. The boundaries separating transport, digital infrastructure and energy systems are becoming increasingly blurred, united by a common dependence on efficient power conversion.

This convergence may ultimately prove to be the conference's defining message. Electrification, artificial intelligence and renewable energy are often discussed as separate industrial trends, yet they are becoming deeply interconnected. Each places unprecedented demands on energy efficiency, power density and system reliability. Each depends upon advances in semiconductor technology. And each requires closer collaboration between disciplines that were once considered distinct.

For the power electronics community, this presents both opportunity and responsibility. The sector is no longer simply enabling technological progress—it is shaping the architecture of the electrified economy itself.

Future leadership will belong not to those who develop the fastest switching device or the most efficient converter in isolation, but to those capable of integrating materials, components, software and manufacturing into complete, intelligent systems.

ANGELTECH made one point abundantly clear: the future of power electronics will be defined less by individual breakthroughs than by the industry's ability to connect them. The systems era has arrived.

Contents

Cover Feature

Rethinking AI data centre cooling

Brooks Henderson of Phononic explain why cooling must be viewed as part of the power delivery and efficiency equation, with thermal management emerging as a core design parameter alongside power and performance.



46

14 Powering the AI mega-rack: how DC-DC conversion is evolving

How the shift to 800 V DC architectures is reshaping power conversion in AI data centres

18 Power electronics at a crossroads?

PE International Conference showcases the technologies and engineering advances driving the next generation of power electronics

23 How Agentic AI is reshaping semiconductor design

Cadence and NVIDIA signal a new phase in electronic design automation, as AI moves to actively orchestrating complex design workflows

28 Advancing SiC substrate and epitaxial technologies

Over the past decade, the SiC power device market has transitioned from a specialty niche to a mainstream enabling technology

32 Innovations in dynamic switch power testing

How Cosmic is redefining reliability testing for wide-bandgap semiconductor devices at production scale

36 Redefining motor drives with 1 MHz GaN switching

James Cannings, CEO at QPT, explores the latest advances in high-frequency GaN motor drive technology and next-generation motor control

42 Rethinking AI power delivery

Hans Hasselby-Andersen, CEO of Lotus Microsystems, explains why conventional power architectures are failing modern AI accelerators



23



60

56 How to pick a hot-swap MOSFET

A comprehensive guide on how to select MOSFETs for server applications and how to evaluate their suitability for hot-swap operation

58 Moore4Power: A new power play for Europe

Infineon is leading a €91 million initiative to advance power electronics, speed development and strengthen Europe's semiconductor resilience

62 Running in the right direction with 48V for AI processor power

Modern AI processors demand unprecedented data throughput, parallel processing performance and continuous multi-kilowatt power delivery

NEWS

06 £10 million Nissan-led project changes the game on EV charging

07 JEDEC releases new SiC guidelines

08 Joint US-Taiwan initiative to scale the future of SiC

09 Polar and Nexperia partner on power MOSFET manufacturing

10 New way to measure WBG device defects

11 RESOLVE- shaping the future of Europe's electronics innovation

12 Infineon opens 'Smart Power Fab' ahead of schedule



06

Editor

Phil Alsop
phil.alsop@angelbc.com

News Editor

Christine Evans-Pughe
christine.evans-pughe@angelbc.com

Technology Editor

Rebecca Pool

Contributing Technical Editor

Richard Stevenson
richard.stevenson@angelbc.com
+44 (0)1923 690215

Sales & Marketing Manager

Shehzad Munshi
shehzad.munshi@angelbc.com
+44 (0)1923 690215

Design & Production Manager

Mitch Gaynor
mitch.gaynor@angelbc.com
+44 (0)1923 690214

Graphic Design & Multimedia Assistant

Harvey Watkins
harvey.watkins@angelbc.com

Sales and Product Manager

James Cheriton
james.cheriton@angelbc.com
+44 (0)2476 718970

Publisher

Jackie Cannon
jackie.cannon@angelbc.com
+44 (0)1923 690205

CEO Sukhi Bhadal

sukhi.bhadal@angelbc.com
+44 (0)2476 718970

CTO Scott Adams

scott.adams@angelbc.com
+44 (0)2476 718970

Power Electronics International Magazine is published six times a year on a controlled circulation basis. Non-qualifying individuals can subscribe at: £105.00/e158 pa (UK & Europe), £138.00 pa (air mail), \$198 pa (USA). Cover price £4.50. All information herein is believed to be correct at time of going to press. The publisher does not accept responsibility for any errors and omissions. The views expressed in this publication are not necessarily those of the publisher.

Every effort has been made to obtain copyright permission for the material contained in this publication. Angel Business Communications Ltd will be happy to acknowledge any copyright oversights in a subsequent issue of the publication. Angel Business Communications Ltd © Copyright 2026. All rights reserved. Contents may not be reproduced in whole or part without the written consent of the publishers. The paper used within this magazine is produced by chain of custody certified manufacturers, guaranteeing sustainable sourcing. © Copyright 2026.

We strive for accuracy in all we publish; readers and contributors are encouraged to contact us if they recognise an error or omission. Once a magazine edition is published [online, in print or both], we do not update previously published articles to align old company names, branding, marketing efforts, taglines, mission statements or other promotional verbiage, images, or logos to newly created or updated names, images, typographic renderings, logos (or similar) when such references/images were accurately stated, rendered or displayed at the time of the original publication.

When companies change their names or the images/text used to represent the company, we invite organizations to provide Angel Business Communications with a news release detailing their new business objectives and/or other changes that could impact how customers/prospects might recognise the company, contact the organisation, or engage with them for future commercial enterprise.

Published by Angel Business Communications Ltd

6 Bow Court,
Fletchworth Gate,
Burnsall Road, Coventry CV5 6SP, UK.
T: +44 (0)2476 718 970
E: info@angelbc.com
W: angelbc.com



£10 million Nissan-led project changes the game on EV charging

UK consortium will use GaN and solar tech to address rising costs of home energy and public charging

ENGINEERS at Nissan Technical Centre Europe (Cranfield, UK) will lead a three year £10 million Government-backed collaborative research project that aims to change how EVs are charged, powered and experienced. Project SUITE - short for Smart Use of Integrated Technology for EVs – addresses the rising costs of home energy and public charging, improving EV efficiency, affordability and sustainability.

The partners are Weeteq (validating AI microcontroller in automotive applications); Oxford PV (developing automotive solar modules); Ricardo (delivering advanced, low-cost, GaN traction inverter); RAM Innovations (developing competency in building high-power GaN modules); University of Bristol (collaborating with industry on semiconductor analysis); Newcastle University (co-creating a digital ecosystem for V2G forecasting); University of Nottingham (advanced design and control of current source inverters); University of Oxford (researching advanced materials for next gen solar cells); and University College London (modelling electricity networks for the optimisation of V2G).

Backed by the UK Government's £4 billion DRIVE35 programme, delivered by the Department for Business and Trade in partnership with the Advanced Propulsion Centre UK (APC) and Innovate UK, the project aims to deliver a number breakthrough technologies.

- Vehicle integrated tandem solar panels to help drivers go further on each charge by topping up the battery with renewable energy and helping reduce energy bills at home.

- An Intelligent Charging Management System (ICMS) to manage customer's electrical consumption and put



renewable energy back into the electrical grid when not in use.

- A bidirectional Automated Charging Device (ACD) for Vehicle to Grid (V2G) connections in real world conditions, improving access to V2G charging and delivering hands-free accessibility of EV charging while maintaining a connection to the grid.

- A high efficiency GaN traction inverter with AI enhanced controls, helping EVs use energy more efficiently, delivering high performance at lower costs to customers.

- Bringing them all together to create an integrated package, combining efficient power electronics, onboard solar, intelligent charging and V2G capability, ready for real-world testing and demonstration. The project is already underway using new 100 percent electric Nissan LEAFs as test vehicles.

David Moss, SVP, R&D, Nissan Africa, Middle East, India, Europe and Oceania (AMIEO) said: "This project highlights the strength of UK innovation and marks a significant step towards the next generation of integrated energy technologies for electric vehicles.

"Through close collaboration with industry and academic partners, we are advancing solutions that lower the cost

of ownership, improve efficiency and deliver greater value for customers."

"From enhanced V2G capability to high performance inverters and solar assisted charging, these developments reinforce the UK's role in building a smarter, more resilient energy ecosystem. Together, we are laying the groundwork for future electrified products and reaffirming our commitment to making electric mobility more accessible, convenient and sustainable for all."

Ian Constance, CEO, Advanced Propulsion Centre UK (APC), said: "Collaborative projects demonstrate the UK's determination to lead the shift to zero-emission mobility. By facilitating the UK Government's DRIVE35 grants, we are turning world-class innovation into industrial capability. With our partners in DBT and Innovate UK, we are backing manufacturers, empowering SMEs, and strengthening the UK's sovereign supply chain.

"This multi-million pound support package is more than an investment in technology; it is an investment in the people, skills, and companies that will define the future of clean transport. Together, we are building the foundations of a competitive, resilient, and sustainable automotive industry."

JEDEC releases new SiC guidelines

Two new documents help advance the reliability, safety, and standardisation of SiC technology

JEDEC Solid State Technology Association has announced the publication of two new documents to support the growing adoption of SiC power semiconductors.

JEP203: Guideline for Short Circuit Evaluation in Power Conversion Transistors, provides guidance for evaluating short-circuit capability of power MOSFETs, helping engineers improve protection design, testing consistency, and system robustness in applications such as EVs, industrial drives, and energy infrastructure.

JEP204: Catalog of Stress Procedures for SiC Devices for Power Electronic Conversion, is a comprehensive reference covering reliability, environmental, and ruggedness stress procedures for SiC power devices.

The publication gives qualification engineers and device manufacturers a common framework for evaluating long-term reliability and performance, helping accelerate industry alignment and increase confidence in next-generation power electronics.

“JEP203 and JEP204 are landmark guidelines for SiC power conversion, enabling the industry to confidently adopt SiC in demanding power electronic applications. JEP204, in particular, provides a comprehensive framework of best practices for stressing SiC power devices, offering significant value to both users and suppliers”, said Donald Gajewski, Wolfspeed senior director reliability engineering, and the chair of the 702_1 Task Group.



“Together, these documents reflect years of industry collaboration and mark a historic milestone in advancing the reliability, safety, and standardisation of SiC technology,”

Thomas Aichinger, senior principal Engineer, SiC technology development, Infineon and the vice-chair of the JC-70.2 subcommittee added: “As the power electronics industry accelerates its transition to SiC, standardisation is the catalyst for confidence and scale.

The release of JEP203 and JEP204 delivers exactly that: clear guidelines for short-circuit protection design and a unified stress-test framework for long-term SiC reliability. These new documents give engineers and manufacturers the tools to design safer systems, validate long-term reliability, and align industry practices.” JC-70

Formed in October 2017 with 22 member companies, JC-70 now has over 70 member companies, underscoring industry commitment to the development of universal standards to help advance the adoption of wide bandgap (WBG) power technologies.

Multinational corporations and technology startups from the US, Europe, Middle East, and Asia are working together to bring to the industry a set of standards for reliability, testing, and parametrics of WBG power semiconductors.

Committee members include industry leaders in power GaN and SiC semiconductors, as well as users of wide bandgap power devices, and test and measurement equipment suppliers. Technical experts from universities and national labs also provide input.

Multinational corporations and technology startups from the US, Europe, Middle East, and Asia are working together to bring to the industry a set of standards for reliability, testing, and parametrics of WBG power semiconductors.

Joint US-Taiwan initiative to scale the future of SiC

Purdue/GCCS initiative targets critical thermal and power bottlenecks in global AI infrastructure

PURDUE UNIVERSITY in the US has signed a five year partnership with Taiwanese company GeChi Compound Semiconductor (GCCS), a specialist in SiC crystal growth, to accelerate the commercialisation of SiC.

The collaboration targets the critical thermal, power and 6G bottlenecks currently constraining the next generation of high-compute infrastructure. Joint research will focus on isolating crystal defects and optimising SiC material growth to accelerate the transition to high-yield 8-inch and 12-inch wafer platforms.

GCCS will serve as a provider of semiconductor materials and Purdue as a hub for the technology. It's the first MOU between Purdue and Taiwan-based GCCS.

"With this agreement, Purdue and GCCS are leveraging research strengths at the academic and industry levels," said Dan DeLaurentis, executive vice president for research. "It continues Purdue's position as a leader in the latest semiconductor research."

"This partnership represents a profound strategic alignment between GeChi Compound Semiconductor and Purdue University," commented Kuan-Ming Hsiung, board chairman. Adding: "By combining our manufacturing scale with America's leading academic institution, we are taking decisive action to secure the domestic supply chain for SiC."



This collaboration is not merely about advancing materials; it is about establishing the resilient, high-yield manufacturing capacity within the United States that is absolutely essential for national tech security and the future of global critical infrastructure."

GE Aerospace qualifies SiC solutions for military vehicles

SiC-based HV controller and converter now production-ready with deliveries to the US Army expected in 2027. GE Aerospace has qualified its High-Voltage Power Controller (HVPC) and Unidirectional Converter (UDC), marking both products as production-ready for US military ground vehicles.

The High-Voltage Power Controller and Unidirectional Converter both use GE Aerospace's SiC MOSFET, Power Overlay Packaging (POL) and module technologies, delivering what are claimed to be best-in-class size, weight, power, and cost advantages.

"The power, temperature and thermal capabilities of these units are unmatched in the industry," said Kris Shepherd, president and general manager, electrical power systems for GE Aerospace. "With this qualification complete, these high-voltage power



units are ready for production and represent a significant technology advancement in power electronics to deliver the power required for military ground vehicles."

Low-rate initial production of the HVPC and UDC is anticipated in 2026, with deliveries to the US Army expected to begin in 2027.

The technology was developed in collaboration with the US Army Combat Capabilities Development Command

Ground Vehicle Systems Center (DEVCOM GVSC) and qualified at GE Aerospace facilities.

These high-voltage units shift from a traditional 28V system to a 600V architecture, enabling more efficient energy conversion and battery charging while reducing weight and easing demand on the motor. This allows the motor to dedicate full power to driving the vehicle when needed.

"This 600V qualification marks a critical milestone in the advancement of high-power vehicle architecture with potential application throughout the US Army's current combat vehicles enabling the next generation of warfighting capability," said Jason Spina, Power System Architect, Ground Vehicle Power & Mobility, Combat Capabilities Development Command (DEVCOM) with the US Army.

Polar and Nexperia partner on power MOSFET manufacturing

Foundry collaboration strengthens supply stability and supports growth in key power markets

POLAR SEMICONDUCTOR and Nexperia have announced a collaboration to make next generation power MOSFET devices at Polar's US based wafer foundry.

By combining Polar's expertise in power semiconductor manufacturing and ongoing capacity expansion with Nexperia's power MOSFET technology, the companies aim to address rapidly growing demand across AI server infrastructure, robotics, industrial, and automotive applications.

Based in Bloomington, Minnesota, Polar provides BCD, BiCMOS, MOSFET, IGBT, MEMS, GaN, and sensor process technologies, along with foundry design enablement, technology transfer, and characterisation services.

Under the collaboration, Nexperia is developing a portfolio of next generation power MOSFETs, spanning a wide range of voltage classes and package types to address diverse power electronics requirements.

"Investing in future technologies and expanding our power MOSFET portfolio are central to Nexperia's strategy," said Fredrik Öberg, general manager business Group MOS at Nexperia.

"Our MOSFETs are recognised for their industry leading quality and reliability, and we are pleased to collaborate with



Polar Semiconductor, a partner that shares our unwavering commitment to zero defect manufacturing and advanced power-device production."

"At Polar, our vision is to create a differentiated wafer foundry focused on power, sensor, and high-voltage semiconductors. Partnering with Nexperia—one of the world's leading power semiconductor companies—to bring their power MOSFET roadmap to market underscores the strength of our strategy and the progress we've made through targeted investment," said Reza Kazerounian, executive chairman of Polar Semiconductor.

Stefan Tilger, interim CEO at Nexperia said: "This collaboration is an important step in strengthening Nexperia's long term manufacturing and supply-chain strategy. By working closely with Polar Semiconductor, we are securing reliable capacity, and providing our customers with a strong, stable, and resilient supply foundation to support future growth."

By combining Polar's expertise in power semiconductor manufacturing and ongoing capacity expansion with Nexperia's power MOSFET technology, the companies aim to address rapidly growing demand across AI server infrastructure, robotics, industrial, and automotive applications.

CGD and NXP to collaborate on application solutions

COMPANIES plan to deliver complete power conversion solutions for automotive, industrial and data centre applications

Cambridge GaN Devices (CGD) is joining forces with NXP Semiconductors to accelerate time to market in data centre and automotive markets.

This new long-term collaboration will enable NXP to develop GaN-based system solutions, building on CGD's GaN products, early access to next-generation CGD GaN developments, and the team's expertise in GaN processes and technologies.

At the same time, CGD will benefit from access to NXP's broad processor and analogue product portfolios, system know-how and global commercial reach, accelerating market penetration through optimised system-level solutions.

Increasing power demands in data centres are particularly driven by AI. A single rack might have only consumed 40 kW in 2022 yet today can draw 200 kW or more and it is expected that a single rack will require 1 MW or more by 2030. CGD says ICeGaN is uniquely positioned to address these requirements enabling higher switching frequency operations, higher power density and delivering superior reliability compared to discrete GaN solutions.

In automotive traction inverters, GaN can improve efficiency at low load. CGD's ICeGaN technology is the only single chip GaN-based technology running on a standard driver that enables paralleling of multiple devices to meet the high current requirements demanded by traction inverters.

New way to measure WBG device defects

US team develops physics-based framework that more accurately detects hidden defects in SiC and GaN

RESEARCHERS at Sandia National Laboratories and Auburn University in the US have developed a new method to more accurately detect atomic-scale defects in electronic materials, an advance that could be particularly useful for wide-bandgap (WBG) semiconductors such as SiC and GaN.

The study 'The completed High-Low method for interface state analysis in MOS capacitors', recently accepted for publication in the *Journal of Applied Physics*, addresses a longstanding challenge in understanding what happens at the critical boundary where a semiconductor meets an insulating layer.

At this interface, microscopic defects can trap electrical charge and quietly reduce device performance, even when the device otherwise appears to function normally. These defects can limit efficiency, increase electrical losses, and reduce the performance of advanced semiconductor devices.

Scientists commonly study these defects by comparing how a device responds to slow and fast electrical signals. However, the technique depends on knowing a key device property, the insulator capacitance, with very high accuracy. Even tiny errors can produce misleading results, sometimes making it appear that far more defects exist than are actually present.

The problem is like tuning a radio. If the frequency is slightly off, the sound becomes distorted. If it is far off, only noise is heard. Only at the correct setting does the signal come through clearly. In these measurements, researchers are effectively trying to "listen" to the defects, with the assumed capacitance acting as the tuning dial.



However, unlike a radio listener who easily recognises a clear broadcast, researchers do not inherently know what the true defect signal should look like. If this tuning parameter is not set exactly, the resulting data cannot be trusted.

Instead of relying on estimates, the researchers developed a physics-based framework that identifies the correct device conditions automatically by enforcing a fundamental electrostatic constraint. In simple terms, all the voltages inside the device must add up consistently according to the laws of physics.

By enforcing this condition, the method removes guesswork and allows researchers to accurately measure defects in regions where the technique previously failed, particularly near the semiconductor band edge, where defects can have the strongest impact on device behaviour.

"This work resolves an otherwise fundamental limitation in one of the most widely used techniques for studying semiconductor interfaces," said Brian D. Rummel, senior member of technical staff at Sandia National Laboratories and lead author of the study. "By introducing a physically consistent framework, we can now extract information that was previously obscured by measurement uncertainty."

The advance is particularly important for SiC and GaN, widely used in high-efficiency and high-power electronics. These materials' performance is often limited by defects at material interfaces. By providing a more accurate way to measure these defects, the new method gives researchers a clearer understanding of what is happening inside the device and how to improve it, according to the team.

"The new analytical framework will help researchers more accurately measure defects in transistor materials," said Sarit Dhar, professor of physics at Auburn University and co-author of the study. "We are excited to apply this method to measure a variety of technologically relevant and exploratory interfaces in our future work."

"Defects at semiconductor interfaces play a major role in determining the performance and reliability of power electronic devices," said Robert J. Kaplar, senior scientist and manager of the semiconductor materials and device physics group at Sandia National Laboratories. "More accurate characterisation tools help us better understand these materials and ultimately improve next-generation electronic technologies."

At its core, the work improves a widely used semiconductor measurement technique by ensuring that it remains fully consistent with fundamental physical principles. The framework helps researchers separate meaningful signals from noise and obtain a clearer picture of the microscopic defects that influence device performance.

As demand grows for faster, more efficient, and more reliable electronics, advances such as this one could help guide the development of next-generation semiconductor technologies.

RESOLVE- shaping the future of Europe's electronics innovation

Eighteen research and technology organizations have committed to RESOLVE, a transnational initiative designed to develop the next generations of electronic components and systems and accelerate their industrial adoption across European territory.

THE INITIATIVE is a concrete response to the sovereignty challenges now posed by quantum technologies and semiconductors in global technological competition.

As major global powers accelerate their investments in disruptive technologies, Europe must transform its assets into industrial and sovereignty levers. RESOLVE addresses this dynamic by proposing an unprecedented framework.

It aims to position Europe as the global leader in future artificial intelligence (AI) solutions, to improve by a factor of 1,000 the energy efficiency of electronic systems by 2032, and to strengthen European semiconductor production capacity, thereby consolidating the continent's technological sovereignty. To achieve this, the project relies on transnational coordination between industry players and 18 research and technology organizations, enabling resource pooling and accelerating the transition from research to market.

Building on the foundations laid by the European Chips Act 1.0 and its pilot lines, this Pan-European initiative reflects the next step Europe now needs: turning strong R&D assets into industrial scale-up, tighter value-chain coordination, and lasting technological sovereignty. RESOLVE embodies a strong vision: making Europe a leader in energy-efficient, secure, and high-performance semiconductor technologies, committed to its industrial competitiveness and strategic independence.

A systemic ambition for Europe at the service of technological mastery and industrial impact

RESOLVE is therefore part of a political ambition for a technological Europe, structured around three priorities:

- **Sovereignty:** developing cutting-edge technologies in key areas, such as advanced memories, power electronics, radio frequencies, photonics, sub-2 nm technologies, and advanced packaging, to reduce dependence on non-European players;
- **Competitiveness:** positioning Europe as a leader in manufacturing semiconductors and high-value-added electronic systems in strategic markets (automotive, aerospace, defense, industry, data centers, etc.);
- **Sustainability:** making energy efficiency a growth driver in support of Europe's climate ambitions to ensure the large-scale deployment of digital infrastructures.

Designed to strengthen European value chains

RESOLVE is structured around two inseparable pillars:

- **"R&D and technology maturation":** 15 key technology areas identified across the semiconductor and electronic systems value chain to

demonstrate feasibility and foster industrial adoption.

- **"From laboratory to industry / From industry to market":** to ensure a rapid and effective transfer to European industry, including early involvement of industrial partners, alignment with product roadmaps, and co-development of prototypes.

The expected outcomes, both technical and economic, will primarily benefit companies in the strategic sectors of the European economy — data centers and AI infrastructure, automotive, defense, space, security and Industry 4.0 — by creating the conditions for new European champions to emerge.

Strengthening Europe's semiconductor industry to sustain its technological leadership

RESOLVE represents a complementary way of building and deploying critical technological capabilities and sovereign industrial foundations in Europe. By bringing together world-unique technological infrastructures into a critical mass, accessible to any European player, RESOLVE weaves a continuum from research to product commercialization, while generating concrete impact on employment, training, and the resilience of value chains, from materials to final systems.

The initiative is thus meant to consolidate and demonstrate Europe's ability to innovate at scale, combining scientific excellence, industrial coordination, and political will. Its ambition is to further strengthen Europe's semiconductor strategy building on proven assets in the areas of energy-efficient, secure, and high-performance semiconductor technologies to guarantee lasting technological leadership and sovereignty.

Building on the foundations laid by the European Chips Act 1.0 and its pilot lines, this Pan-European initiative reflects the next step Europe now needs: turning strong R&D assets into industrial scale-up, tighter value-chain coordination, and lasting technological sovereignty

Infiniteon opens 'Smart Power Fab' ahead of schedule

Investment of €5 billion strengthens manufacturing base for modern power semiconductors, creating 1,000 jobs

INFINEON has opened the Smart Power Fab in Dresden, Germany several months ahead of schedule.

The €5 billion plant is the largest single investment in Infineon's history and one of the largest investment projects in Germany, creating 1,000 new, direct jobs.

The new fab doubles Infineon's manufacturing capacities at the Dresden site, thus creating the world's largest factory for intelligent power semiconductors and analogue / mixed-signal technologies.

"We're opening our new plant at just the right time," says Jochen Hanebeck, CEO of Infineon Technologies AG. "Our Smart Power Fab is creating urgently needed capacities for the key technologies of the future, for everything from energy supply for AI data centres to software-defined vehicles and renewable energies. Infineon is thus giving an important impulse in making the global AI revolution possible and securing supply chains in critical industries. By taking this step, we are strengthening our global vanguard position as a leading manufacturer of power semiconductors and analog / mixed-signal technologies."

"Infineon's new Smart Power Fab sends out a powerful message about Germany and Europe as a centre of industry," says German Federal Chancellor Friedrich Merz. "This investment shows that competitive, cutting-edge semiconductor production in our country is possible. Power semiconductors are a key technology for the energy transition, the future of mobility and our AI infrastructure – meaning they are crucial to make our economy fit for the future. The expansion of Infineon's production capacity in Dresden also strengthens



our technological sovereignty and the resilience of important supply chains in Europe."

Karsten Wildberger, German Federal Minister for Digital Transformation and Government Modernization, says: "What is being produced here in Dresden secures jobs and creates value in Germany, because these cutting-edge chips are at the heart of key technologies such as electric vehicles, wind turbines, and the power infrastructure for AI data centres. The Smart Power Fab marks another milestone on Germany's path toward digital sovereignty and demonstrates that Germany can move fast, Germany can deliver high-tech innovation, and Germany can shape the future."

Saxony's Minister-President Michael Kretschmer says: "Microelectronics is the key technology of the 21st century, the foundation of innovation, competitive strength and economic progress. The choice of Dresden, home to Europe's largest microelectronics cluster, represents a decisive location advantage. Infineon's new Smart Power Fab makes this unique ecosystem stronger and shows that major industrial projects can be implemented securely and swiftly with the right basic

prerequisites. This sends a strong signal about Dresden as an industrial location and about the future of microelectronics in Europe."


The new Smart Power Fab uses digitalisation to accelerate production to twice as fast as in the past, depending on demand. For example, the building itself and the ideal machine layout were pre-planned and optimised using a digital twin. System and process clearance is supported by AI algorithms. Furthermore, the connection to the Infineon plant in Villach as "One Virtual Fab" makes qualification of processes and products significantly faster than in the past. The increased ramp-up speed lets Infineon effectively and flexibly leverage current market opportunities, including in the field of AI.

Infineon says its Smart Power Fab is also setting new standards in sustainable manufacturing. The new manufacturing system requires no natural gas and cuts water consumption with water treatment and the implementation of closed-loop water systems. Approximately 90 percent of the water used is to be reintroduced in the circulation, making it possible to recover as much as 45 percent of the energy used.




www.EVGroup.com

HIGH-VACUUM BONDING



EVG® ComBond®

Oxide-free, covalent wafer bonding
High-vacuum process environment
Room-temperature process



GET IN TOUCH to discuss your manufacturing needs
www.EVGroup.com

Powering the AI mega-rack: how DC-DC conversion is evolving for AI data centres

Advanced Energy Director Harry Soin tells Power Electronics how the shift to 800 V DC architectures is reshaping power conversion in AI data centres.

WITH THREE GENERATIONS of AI computing platforms in as many years, NVIDIA has rapidly increased compute density while driving a transition toward rack-scale 800 V DC power architectures.

At the same time, hyperscalers including Google, Meta, and Microsoft are collaborating on Open Compute Project (OCP) initiatives, such as Mount Diablo, which re-architect data centre power distribution to provide a pathway to megawatt-scale AI racks at higher voltages. The impact on power conversion has been profound with industry players developing systems set to provide higher-voltage DC-DC conversion and greater

efficiencies at megawatt-scale rack power levels.

A case in point is the ADH series of DC-DC converters from US-based Advanced Energy, introduced in June this year. Targeting 800 V DC AI data centre power architectures, these DC-DC converters can convert that 800 V DC input into 50 V DC output and operate at multi-megahertz switching frequencies. Devices are said to achieve a peak power efficiency of 98.2%, helping to reduce power losses in high performance computing applications. Each converter can deliver up to 8 kW of peak power and 6 kW of continuous full-load power within a standard half-brick package.

“From the Grace Hopper to Vera Rubin Ultra NVIDIA platforms, power is essentially doubling and quadrupling, so currents are increasing by huge amounts and power dissipation losses are significantly higher,” highlights Harry Soin, Senior Director, Hyperscale, at Advanced Energy. “To keep up with this, we transitioned from 12 V to 48 V outputs... and with industry now moving towards 800 V architectures, we have introduced our latest converters.”

Megawatt power racks

The transition to 800 V DC has been rapid. In 2022 standard power shelves from Advanced Energy handled 18 kW but within that same year, capacity swiftly rose to 33 kW. This year, the company will be releasing a 72 kW power shelf, which can be either combined or used in an external sidecar configuration to reach the 800 kW to full megawatt power racks data centres are now demanding.

But once megawatt scale is well and truly reached, traditional power shelves hit a bottleneck; distributing this much power through standard low-voltage copper busbars requires them to become prohibitively thick to avoid over heating. Given this, next-generation power rack architectures, such as NVIDIA's Kyber platform, designed to scale to the massive NVL576 architecture with 576 connected Rubin Ultra GPUs, removes the power shelf from the main cabinet and instead runs a 800V DC power line directly into the server.

“GPU suppliers are working on a solution where essentially, the high



➤ Keeping cool: a Google data centre aisle, with coolant distribution units, at New Albany, Ohio. [Google]



➤ The NVIDIA Rubin platform: The AI supercomputer architecture is designed to power massive agentic AI. [NVIDIA]

voltage goes directly to the server - this is a revolutionary move, right?" says Soin. "This means the server itself needs power conversion."

According to Soin, Advanced Energy's DC-DC team is currently working on Power Distribution Boards (PDBs) to do just this. These will house 4 kW to 6 kW DC-DC converters in a standard brick package – and four of these bricks can be placed onto the server's PDB. This will then step down the incoming 800 V to low voltages at the board level, to power the GPUs, such as Rubin Ultra.

As the demand for high power density AI infrastructure - that generates hugely concentrated thermal loads - continues to grow, the move from air cooling to liquid cooling will be critical for AI workloads. As Soin points out, he and colleagues have been focused on developing DC-DC converters for the various GPU architectures that link tens and tens of GPUs together.

However, the architectures that accompany the Vera Rubin and upcoming Kyber platforms double the computing density of existing systems. This extreme density is necessary to power the massive agentic AI models that process gargantuan levels of data.

"We're now working with less space for the power and we need for these high-density platforms to be more efficient... If you can't get the heat out, you cannot provide the power, so managing heat with direct liquid cooling has become key, especially for architectures like the Kyber platform," says Soin.

With this in mind, Advanced Energy has designed its board-mounted DC-DC converters so that liquid cooling can be directly integrated via a baseplate.

"By tapping into the liquid plumbing used by all the hyperscalers and chip suppliers, including AMD and NVIDIA, we can directly remove the heat," explains Soin.

"We used to be working with 1 to 2 kW in a single brick board mount power module but are now working with 6 to 8 kW modules and are considering 12 kW modules – this is only possible with direct liquid cooling," he adds.

Cooling aside, Soin and colleagues at Advanced Energy are working with both silicon and wide bandgap materials to take the peak power efficiency of their devices beyond 98.2%. Soin says they have embraced silicon carbide semiconductors, thanks to solid reliability reports, and are now using gallium nitride as more and more reliability data becomes available.

"We're now working with less space for the power and we need for these high-density platforms to be more efficient... If you can't get the heat out, you cannot provide the power, so managing heat with direct liquid cooling has become key, especially for architectures like the Kyber platform"

"We are looking at incremental [increases] in efficiency ... and the only way to move forward and reach [these] now is to move beyond silicon," he says. "We're removing the parasitics and the switching losses. After that, it's going to be a race for the conduction losses and this can only be achieved with the lower on-resistance found in wide bandgap devices."

Market moves

The Vera Rubin Ultra platform with its multi-chip GPU design is, according to NVIDIA, expected to launch in the second half of 2027, at which point Advanced Energy anticipates rising demand for its ADH series of DC-DC converters. As Soin puts it: "We're developing solutions for hyperscalers as well as the GPU designers — these guys are leading this revolution, right?"

"We are designing solutions and reference designs that could be part of their ecosystem and will essentially be deployed," he adds. "Public information indicates that Vera Rubin Ultra will be launched in huge quantities in the middle of next year, and that's exactly when you will see our [devices] verified and implemented — they will be needed to address the power needs of those upcoming architectures."

As today's data centres transition to 800 V DC via additional sidecars to meet rising AI energy demands, Soin reckons native 800 V infrastructure will come online around mid-2029. "When that 800 V goes directly to the server and [delivers] higher efficiencies – that really will shape the future."

Soin is also clear that industry will not stop at 800 V. "If power keeps increasing as it is today, then we could approach 1500 V DC in the future," he says. "This is a really exciting time."

20
27

Sheraton Brussels
Airport Hotel

Brussels, Belgium

12-14 April 2027



REGISTER NOW

Power Electronics International Conference 2027

Building on four hugely successful events, PE International returns for its 5th year. Chaired by Aleksei Cherkasov from Infineon, the conference brings together executives, end users, and technology experts for two days of insight, discussion, and collaboration on the technologies driving the global power electronics industry.

The 2027 programme builds on the core themes from 2026; covering automotive, renewables, data centres, and wide bandgap innovation, while placing a stronger emphasis on roadmaps and bridging the gap between research and manufacturing.

peinternational.net



SCAN ME

80+ EXHIBITORS

800+ ATTENDEES

4 CO-LOCATED CONFERENCES

120+ PRESENTATIONS





The European Event for The Semiconductor Supply Chain

REGISTER NOW

Conference Themes

- AI Data Centres from Grid to Rack
- Electrification at Scale: EVs, Charging and High Voltage
- Wide Bandgap Dominance
- Power Semiconductors in Action: Application Focus

Why Attend?

- Unbeatable location: The Sheraton Brussels Airport Hotel is 20 yards from the airport terminal — fly in, walk to the conference.
- Networking: Two drinks receptions and a conference dinner provide exceptional networking opportunities.
- Full supply chain: Four co-located conferences bring the entire compound semiconductor, photonics, power electronics, and advanced packaging supply chain together under one roof.
- 800+ attendees: Join a growing community of decision-makers, engineers, and researchers.

THE WHOLE SEMICONDUCTOR SUPPLY CHAIN UNDER ONE ROOF



PREVIOUS SPEAKERS & SPONSORS





Power electronics at a crossroads?

The PE International Conference, part of the highly successful ANGELTECH event held earlier this year, highlighted the technologies shaping the power electronics industry. From automotive electrification and AI-driven data centre densification to engineering innovation with wide band gap (WBG) devices and power semiconductors for the future grid, the programme painted a clear picture of an industry undergoing rapid systems-level transformation.

THE AUTOMOTIVE sessions demonstrated how electrification is reshaping vehicle architecture around power electronics, with SiC and GaN technologies enabling higher efficiency, faster charging and more intelligent energy management. The AI infrastructure discussions revealed that data centres are entering a new era of extreme power density, where cooling, energy delivery, batteries and intelligent infrastructure must operate as fully integrated systems to support GPU-driven workloads.

Meanwhile, the WBG-focused presentations highlighted how silicon carbide and gallium nitride are moving from emerging technologies into mainstream deployment across automotive, industrial, renewable energy and digital infrastructure applications. The power grid theme reinforced how advanced

semiconductors are becoming critical to future energy infrastructure, supporting renewable integration, intelligent grid management, HVDC systems and more efficient power conversion.

Collectively, the conference reinforced that future competitiveness across multiple industries will depend on the ability to combine advanced semiconductors, thermal management, power conversion, packaging, manufacturing scalability and intelligent control into highly efficient, resilient and scalable systems.

Putting the automotive industry in pole position

At PE International's "Putting the Automotive Industry in Pole Position" theme, one message came through clearly: the automotive sector is now one of the most important battlegrounds for the future of power

electronics. Across the presentations, speakers explored how electrification, energy efficiency, supply chain resilience, and semiconductor innovation are reshaping the global automotive landscape.

The theme focused heavily on the transition from traditional internal combustion engines toward electrified mobility platforms. While battery electric vehicles remain central to long-term decarbonisation goals, the discussions recognised that the market is becoming more nuanced. Hybrid systems, plug-in hybrids, and advanced powertrain architectures are all playing a role as OEMs balance cost, infrastructure readiness, consumer demand, and regulatory pressure. The result is a rapidly diversifying automotive ecosystem where power electronics sit at the centre of vehicle performance, efficiency, and reliability.

A major theme throughout the event was the growing strategic importance of wide band gap semiconductor materials such as silicon carbide (SiC) and gallium nitride (GaN). These technologies are increasingly enabling higher switching frequencies, reduced thermal losses, improved efficiency, and more compact power conversion systems. For automotive manufacturers, this translates directly into longer driving range, faster charging, reduced system weight, and more efficient energy management.

The presentations also highlighted how the vehicle is evolving into a highly integrated electrical architecture. Modern EV platforms are no longer simply mechanical systems with electronic add-ons. Instead, they are becoming software-defined, power-managed systems requiring sophisticated integration between semiconductors, modules, thermal management, sensing, battery systems, and digital controls. This convergence is driving collaboration across the semiconductor, packaging, automotive, and energy sectors.

Another significant discovery was the extent to which automotive electrification is influencing the broader industrial ecosystem. Supply chain resilience emerged as a recurring concern, particularly around critical materials, semiconductor manufacturing capacity, and battery production.

Europe's automotive industry faces growing pressure from global competition while simultaneously trying to accelerate electrification targets and localise manufacturing capabilities. Discussions reflected wider industry concerns that competitiveness, sustainability, and industrial policy are now inseparable.

Thermal management and power density were also key technical themes. As EV architectures move toward higher voltage systems, including 800V platforms, engineers are being challenged to manage increasingly demanding thermal loads while maintaining reliability and safety. The conference explored how advanced packaging, cooling technologies, and module integration are becoming critical enablers of next-generation vehicle platforms.

Importantly, the event reinforced that automotive innovation is no longer isolated within OEMs. Utilities, charging infrastructure providers, semiconductor manufacturers, materials companies, and software developers are now part of the same interconnected ecosystem. The automotive sector is becoming a systems-level industry where energy infrastructure, grid stability, battery technology, and intelligent power management all intersect.

Ultimately, "Putting the Automotive Industry in Pole Position" positioned power electronics as the enabling foundation of the automotive transition. The conference demonstrated that future vehicle leadership will not be defined solely by mechanical engineering or manufacturing scale, but by the ability to integrate advanced semiconductors, intelligent energy systems, and resilient supply chains into highly efficient electrified platforms. In that sense, the event captured a broader industrial reality: the race for automotive leadership is increasingly a race for power electronics leadership.

Delivering AI-driven data centre densification


Artificial intelligence is fundamentally reshaping the data centre industry, and the "Delivering AI-driven Data Centre Densification" theme at PE International highlighted just how quickly infrastructure requirements are changing. Across the presentations, one message became clear: AI is not simply increasing compute demand. It is redefining power architecture, thermal management, rack density, energy

strategy, and the entire operational model of the modern data centre.

The conference explored how AI workloads are driving an unprecedented increase in power density. Traditional enterprise data centres were historically designed around racks operating at 5–10kW. AI clusters, however, are now pushing well beyond 30kW, 50kW and in some cases exceeding 100kW per rack. This shift is forcing operators, infrastructure providers and technology vendors to rethink almost every aspect of facility design.

One of the key lessons from the presentations was that densification is no longer purely about fitting more compute into the same physical footprint. Instead, it has become an exercise in balancing power delivery, thermal stability, energy efficiency and operational resilience simultaneously. AI infrastructure introduces highly dynamic load profiles that differ significantly from traditional workloads. GPU clusters create rapid and unpredictable power spikes, placing new stress on UPS systems, batteries, cooling infrastructure and grid connectivity.

As a result, power electronics are becoming central to the future of AI infrastructure. Multiple presentations highlighted how the sector is moving toward more sophisticated power conversion architectures capable of delivering higher efficiency and faster transient response. Wide band gap semiconductor technologies such as silicon carbide and gallium nitride featured prominently because of their



ability to improve switching efficiency, reduce losses and support higher power densities within more compact systems.

Thermal management emerged as one of the defining technical challenges of AI-driven densification. Traditional air-cooling approaches are increasingly struggling to support the thermal loads generated by advanced AI accelerators.

Several presentations focused on the transition toward liquid cooling technologies, including direct-to-chip cooling and immersion cooling strategies. These approaches are rapidly moving from niche deployments into mainstream hyperscale planning.

What became particularly apparent was that cooling is no longer an isolated mechanical issue. Instead, it is becoming deeply integrated with power strategy and overall system architecture. Higher rack densities require closer coordination between cooling systems, electrical infrastructure and building design. Operators are now being forced to think holistically about how power and heat interact across the entire facility.

Energy availability and grid interaction also featured heavily throughout the conference. AI-driven data centre growth is occurring at a time when many regions are already facing grid constraints and rising electricity demand. Presenters repeatedly stressed that access to reliable power is becoming one of the most important determinants of future data centre deployment. This is creating growing interest in alternative energy strategies, microgrids, on-site generation and battery-backed resilience models. Several speakers highlighted how data centres are increasingly behaving like active participants within wider energy ecosystems rather than passive consumers of electricity.

Intelligent energy management, load balancing and real-time optimisation are becoming critical operational capabilities.

Another major lesson from the event was the increasing convergence between digital infrastructure and industrial engineering. Delivering AI-scale facilities now requires deep collaboration across power engineering, semiconductor technology, cooling, software management, construction and energy planning. The traditional boundaries between IT infrastructure and industrial infrastructure are rapidly disappearing.

The presentations also explored how AI itself is being used to optimise data centre operations. AI-driven monitoring systems are increasingly capable of predicting thermal hotspots, managing energy efficiency, optimising workloads and improving preventative maintenance. This creates a feedback loop where AI infrastructure not only consumes vast amounts of compute power but also becomes essential to managing the complexity of that infrastructure efficiently.

Importantly, the conference highlighted the growing role of advanced batteries and energy storage technologies in supporting AI workloads. Traditional backup power assumptions are evolving. Instead of simply providing emergency power during outages, modern battery systems are increasingly expected to respond dynamically to rapid load fluctuations created by AI processing demands. Fast-response power delivery is becoming just as important as resilience. Supply chain resilience and manufacturing scalability also emerged as major concerns. AI infrastructure growth is placing enormous pressure on semiconductor supply chains, power component manufacturing and cooling technology deployment. Several

presentations reflected broader industry concerns around how quickly the ecosystem can scale to meet demand while maintaining reliability, efficiency and sustainability.

Sustainability itself was another recurring theme. AI workloads are dramatically increasing energy consumption, placing pressure on operators to improve efficiency metrics while continuing to scale compute capacity. This is accelerating interest in higher-voltage distribution architectures, advanced power management systems and more efficient cooling technologies.

The discussions also revealed a significant change in data centre design philosophy. Historically, facilities were often optimised around redundancy and uptime. While resilience remains essential, AI infrastructure is increasingly demanding optimisation around power efficiency, scalability and workload flexibility. Operators are now being forced to design for continuous evolution as AI hardware, cooling requirements and energy demands change at unprecedented speed.

One particularly important takeaway was that AI-driven densification is not solely a hyperscaler issue. Enterprise operators, colocation providers and edge infrastructure companies are all beginning to face similar challenges as AI workloads expand beyond centralised cloud environments. This means the entire industry must adapt to a future where high-density compute becomes increasingly normalised.

The conference also reinforced the strategic importance of infrastructure planning at a regional and national level. Data centres are now central components of economic competitiveness, AI capability and digital sovereignty. Governments,

Artificial intelligence is fundamentally reshaping the data centre industry, and the “Delivering AI-driven Data Centre Densification” theme at PE International highlighted just how quickly infrastructure requirements are changing. Across the presentations, one message became clear: AI is not simply increasing compute demand. It is redefining power architecture, thermal management, rack density, energy strategy, and the entire operational model of the modern data centre.

utilities and infrastructure providers are becoming more deeply involved in questions around energy access, planning approvals and long-term infrastructure investment.

Ultimately, the “Delivering AI-driven Data Centre Densification” theme demonstrated that the future of data centres will be defined by integrated systems thinking. Success will depend on the ability to coordinate power electronics, cooling, energy management, semiconductors, batteries, software intelligence and grid interaction as part of a unified infrastructure platform.


The presentations made clear that AI is no longer simply another workload category. It is a structural force transforming the physical and operational foundations of digital infrastructure itself. As rack densities rise, energy demands increase and infrastructure complexity grows, the industry is entering a new era where efficiency, resilience and intelligent power management will determine competitive advantage.

Engineering innovation with WBG devices, circuits and materials

The “Engineering Innovation with WBG Devices, Circuits and Materials” theme at PE International highlighted how wide band gap (WBG) technologies are rapidly moving from emerging innovation to mainstream industrial deployment. Across the presentations, the conference explored how silicon carbide (SiC) and gallium nitride (GaN) are transforming power electronics at device, circuit and systems level, enabling a new generation of higher efficiency, higher power density and higher performance applications.

One of the clearest lessons from the presentations was that WBG materials are no longer viewed simply as replacements for silicon. Instead, they are becoming enabling technologies that fundamentally change how power systems are designed. Their ability to operate at higher voltages, higher temperatures and higher switching frequencies allows engineers to rethink converter architectures, thermal strategies and overall system integration.

The conference repeatedly demonstrated that the adoption of WBG



technologies is accelerating because traditional silicon is approaching practical performance limits in many applications. Electrification trends across automotive, renewable energy, industrial automation, aerospace and data centres are demanding greater efficiency and increased power density simultaneously. WBG devices are increasingly viewed as essential to meeting those requirements.

Silicon carbide emerged as one of the dominant themes throughout the event. Multiple presentations explored its growing role in electric vehicle traction inverters, high-voltage power conversion and fast-charging infrastructure. SiC's lower switching losses and higher thermal conductivity enable smaller passive components, reduced cooling requirements and improved overall system efficiency. For automotive applications, this translates directly into increased driving range, faster charging capability and reduced system weight.

The conference also highlighted how SiC is reshaping high-power industrial applications beyond automotive. Renewable energy systems, grid infrastructure and industrial motor drives are all benefiting from the efficiency improvements made possible by higher switching frequencies and reduced energy losses. As global electrification accelerates, these efficiency gains become increasingly important at system level.

Gallium nitride was another major focus, particularly in applications where high-frequency switching and compact system design are critical. Presenters demonstrated how GaN devices are enabling smaller, lighter and faster power conversion systems, particularly

in consumer electronics, telecoms and emerging AI infrastructure. GaN's ability to operate efficiently at extremely high switching frequencies opens opportunities for miniaturisation that are difficult to achieve with silicon or even silicon carbide in certain applications.

A recurring lesson from the conference was that the real innovation is no longer occurring solely at the device level. Instead, the greatest engineering challenges and opportunities are now emerging around system integration. WBG devices require new approaches to gate driving, electromagnetic interference management, thermal engineering, packaging and circuit design. Several presentations emphasised that simply replacing silicon devices with SiC or GaN equivalents does not automatically deliver optimal performance. Engineers must redesign converter topologies, control strategies and packaging approaches to fully exploit the capabilities of WBG materials. This systems-level engineering challenge is becoming one of the defining themes of modern power electronics.

Advanced packaging technologies also featured prominently throughout the discussions. As switching speeds increase and power densities rise, packaging increasingly determines overall device reliability and performance. Presenters explored innovations in module integration, interconnect technologies, thermal interfaces and substrate materials designed to support next-generation WBG systems.

Thermal management was another critical topic. Although WBG devices operate more efficiently than silicon,

the increased power densities they enable create new thermal engineering challenges. Several presentations examined how advanced cooling strategies, materials engineering and package design are becoming essential for maintaining reliability in high-performance systems.

Reliability and manufacturability were recurring concerns across the conference. While WBG devices offer major performance advantages, large-scale adoption requires consistent manufacturing quality, predictable long-term behaviour and scalable production capacity. Discussions reflected the industry's growing focus on defect reduction, wafer quality improvement and process maturity as SiC and GaN move into mass-market applications.



Supply chain resilience also emerged as an important issue. As demand for WBG devices accelerates globally, manufacturers are facing growing pressure around substrate availability, epitaxial growth capacity and fabrication scalability. Several presentations highlighted the strategic importance of securing regional semiconductor manufacturing capabilities, particularly within Europe and North America.

One particularly important lesson from the event was the increasing convergence between materials science and systems engineering. Progress in WBG technologies is no longer driven solely by semiconductor physics. Instead, innovation now depends on coordination across crystal growth, epitaxy, device fabrication, packaging, thermal engineering, circuit design and software control.

The conference also demonstrated how WBG technologies are enabling broader trends in electrification and sustainability. Improved efficiency at device level translates into substantial energy savings when deployed across large-scale industrial systems, EV fleets

or energy infrastructure. As global electricity demand continues to rise, these incremental efficiency gains become strategically significant.

Artificial intelligence and data centre infrastructure also appeared as emerging application areas for WBG technologies. Higher efficiency power conversion is becoming increasingly critical as AI workloads drive rapid increases in data centre power consumption. WBG devices are enabling more compact and efficient power delivery systems capable of supporting these next-generation compute environments.

Importantly, the presentations reinforced that WBG adoption is now transitioning from early-stage

experimentation into commercial scale deployment. Automotive platforms, renewable energy systems and industrial power architectures are increasingly being designed around SiC and GaN from the outset rather than treating them as optional upgrades.

The conference also highlighted the growing importance of collaboration across the ecosystem. Device manufacturers, materials suppliers, packaging specialists, OEMs and research institutions are all contributing to accelerating WBG deployment.

The pace of innovation is increasingly being driven by ecosystem-level integration rather than isolated technology development. Ultimately, "Engineering Innovation with WBG Devices, Circuits and Materials" demonstrated that wide band gap technologies are becoming foundational to the future of power electronics. The event showed that SiC and GaN are not simply incremental improvements over silicon, but technologies capable of reshaping how energy is converted, managed and distributed across multiple industries.

As electrification accelerates globally, WBG materials are emerging as critical enablers of higher efficiency, greater power density and more intelligent energy systems. The conference made clear that future leadership in power electronics will depend not only on advances in semiconductor devices themselves, but on the ability to integrate materials, circuits, packaging and systems engineering into highly optimised next-generation platforms.

Ensuring a renewable, sustainable future

This conference theme explored how next-generation power semiconductors will play a critical role in enabling the future electrical grid. As electrification accelerates across transport, industry, renewable energy and AI infrastructure, grid systems are being placed under increasing pressure to deliver greater efficiency, resilience and flexibility.

The presentation highlighted how wide band gap technologies such as silicon carbide (SiC) and gallium nitride (GaN) are becoming essential for modern grid infrastructure. Their ability to operate at higher voltages, temperatures and switching frequencies enables more efficient power conversion, reduced energy losses and more compact system designs. These capabilities are increasingly important as utilities integrate renewable energy sources, battery storage and distributed power generation into existing networks.

A key lesson from the session was that future grids will require far more intelligent and dynamic power management. Advanced semiconductors are enabling faster switching, improved control systems and higher power density across applications such as HVDC transmission, renewable energy integration, fast charging infrastructure and smart substations. The presentation also reinforced the importance of manufacturing scalability and supply chain resilience. As global electricity demand rises, the ability to produce reliable, high-performance semiconductor technologies at industrial scale will become increasingly important.

Ultimately, the session demonstrated that power semiconductors are becoming foundational technologies for delivering a more efficient, flexible and electrified energy future.



How Agentic AI is reshaping semiconductor design

The expanded collaboration between Cadence and NVIDIA signals a new phase in electronic design automation, where AI moves beyond assisting engineers to actively orchestrating complex design workflows. As semiconductor complexity accelerates, agentic AI, digital twins and accelerated computing are converging to redefine how chips, systems and entire AI factories are conceived, verified and optimised. This article is based on an interview with Rob Knoth, Sr Solution Marketing Group Director, Cadence.

THE SEMICONDUCTOR industry has always been defined by relentless escalation. Every new generation of processors, networking devices, accelerators and systems presents challenges that could not be solved using the methodologies of the previous generation. From shrinking geometries and rising transistor counts to advanced packaging and heterogeneous integration, engineers have consistently been forced to rethink not only what they build, but how they build it.

Today, that pattern is repeating itself on a much larger scale.

Artificial intelligence has become both the driving force behind next-generation semiconductor infrastructure and the technology transforming the tools used to create it. This dual role is central to how Cadence views the future of engineering: “design for AI” and “AI for design.”

The first concerns enabling customers such as NVIDIA to create increasingly

sophisticated AI infrastructure. The second may prove even more transformational, applying the latest AI technologies directly to engineering workflows themselves.

The result is a fundamental shift from traditional automation towards autonomous, agent-driven design environments that promise to deliver unprecedented productivity gains while freeing engineers to focus on innovation rather than administration.

Engineering's Next Inflection Point

For engineers with decades of experience, the current AI revolution feels both familiar and unprecedented. The semiconductor industry has always evolved through successive waves of technological advancement.

With each generation, engineers have required better tools, faster computation, more sophisticated simulation capabilities and improved collaboration methods to tackle problems that were previously unsolvable.

According to Cadence, what distinguishes the current moment is the extraordinary pace at which AI itself is advancing.

Since the emergence of large language models and generative AI, the technology landscape has evolved at a speed rarely seen in engineering disciplines. Organisations that fail to continuously reassess AI capabilities risk falling significantly behind what is achievable.

This reality has intensified collaboration between EDA vendors and semiconductor leaders. Cadence's relationship with NVIDIA exemplifies this dynamic. On one side, Cadence provides the design software, IP, hardware platforms and verification infrastructure that enable NVIDIA to develop successive generations of AI accelerators and systems. On the other,

Cadence is incorporating NVIDIA-powered AI technologies directly into its own design platforms.

What emerges is a virtuous cycle: AI infrastructure enables increasingly powerful AI models, which in turn help engineers create the next generation of AI infrastructure.

Beyond Traditional EDA Automation

Automation has always been at the heart of Electronic Design Automation. Indeed, the "A" in EDA has historically referred to simplifying repetitive engineering tasks.

Traditional EDA automation, however, relied heavily on deterministic workflows and rule-based logic. Engineers encoded decision trees, heuristics and predefined actions that software could execute automatically. This approach remains valuable, but its limitations become increasingly apparent as design complexity grows. Agentic AI introduces a fundamentally different model.

Instead of simply executing predefined instructions, AI agents can interpret goals, evaluate outcomes, determine next steps and orchestrate complex workflows with minimal human intervention.

The distinction is significant. In earlier AI-enhanced workflows, verification tools might automatically launch simulations, generate

reports and return log files for analysis. Engineers would still spend considerable time interpreting results, diagnosing failures and determining subsequent actions.

Agentic systems aim to automate much of that loop.

Cadence recently demonstrated an engineering proof of concept developed jointly with NVIDIA in which agentic AI orchestrated a verification workflow around NVIDIA intellectual property. According to the companies, the result delivered a dramatic reduction in runtime, shrinking processes that previously required weeks into workflows measured in days.

While these systems are not yet replacing human engineers, they are demonstrating a new level of workflow autonomy that was previously unattainable.

The significance lies not merely in speed, but in the broader implications for engineering productivity.

The Rise of the Autonomous Engineer

The concept of an "autonomous engineer" may sound futuristic, but industry leaders are careful to distinguish vision from reality.

Fully autonomous chip development remains a distant objective.

Semiconductor design encompasses countless decisions involving architecture, trade-offs, market requirements and innovation that continue to require human expertise. However, the ability to automate substantial portions of verification, implementation and optimisation workflows is already becoming practical.

Cadence's evolving AI strategy includes platforms such as ChipStack, Verisium and Innovus-related agentic workflows that combine AI agents with traditional EDA environments to produce trusted, verifiable results.

Rather than replacing engineers, these systems function as highly capable collaborators.

This distinction is crucial.



Engineering has traditionally involved significant amounts of “nanny work”: monitoring jobs, restarting failed processes, analysing logs, debugging tool issues and managing computational workflows. Many highly skilled engineers spend substantial portions of their careers performing tasks that contribute little to creative innovation.

Agentic AI targets that burden directly. By eliminating routine monitoring and administrative effort, engineers can devote greater attention to architecture, creative problem solving and system-level innovation.

For an industry facing persistent talent shortages and escalating design complexity, that shift could prove transformative.

Unlocking Human Creativity

The productivity gains associated with agentic AI extend beyond reducing drudgery.

Perhaps the larger opportunity lies in expanding what engineers can accomplish.


Semiconductor systems increasingly involve multiple domains simultaneously, including electronics, software, thermal management, packaging, optics and communications. As these domains converge, identifying relationships and optimisation opportunities becomes more difficult for individuals working within traditional design silos.

Agentic AI offers a mechanism for bridging those gaps.

By integrating multiple tools, datasets and simulation environments, AI agents can help engineers explore interactions that might otherwise remain hidden.

This capability could accelerate breakthroughs in advanced packaging, heterogeneous integration and multi-physics optimisation.

Instead of spending valuable time managing tools, experienced engineers may increasingly focus on discovering new architectural concepts and solving previously inaccessible problems. Such a shift could have important implications for attracting and retaining engineering talent.



Highly creative individuals are often discouraged by repetitive workflow management. Reducing that burden may help ensure that engineering careers remain intellectually rewarding while enabling experienced innovators to remain engaged longer.

Digital Twins Become Strategic Assets

As AI infrastructure scales, digital twins are emerging as a critical component of engineering strategy.

Historically, digital twin technology has been associated with manufacturing and industrial systems. However, the concept is rapidly expanding into semiconductor development and data centre optimisation.

Cadence argues that the methodology underpinning semiconductor design already demonstrates the value of simulation-first development. More than 99% of a semiconductor’s behaviour is typically simulated before fabrication. This allows engineers to identify and correct issues long before expensive manufacturing processes begin.

By comparison, many large-scale physical systems rely on much lower levels of simulation prior to physical testing.

Data centres represent a particularly promising application area.

Unlike semiconductor devices, which remain fixed after fabrication, data centres are dynamic environments.

Equipment is continuously upgraded, workloads are redistributed, cooling systems are adjusted and power delivery architectures evolve over time. A high-fidelity digital twin allows engineers and operators to model these changes before implementing them physically.

Combined with accelerated computing and AI-based optimisation, digital twins can improve power utilisation, predict performance bottlenecks and support more efficient infrastructure expansion. As AI deployments continue to scale globally, these capabilities are becoming increasingly important.

From PUE to Tokens per Watt

The rapid growth of AI infrastructure is also prompting a reassessment of efficiency metrics.

For years, Power Usage Effectiveness (PUE) has served as the standard measure of data centre efficiency. While valuable, PUE primarily measures how effectively a facility converts incoming energy into useful computing power.

Many industry observers now argue that PUE alone does not fully capture the value generated by AI systems. This has contributed to growing interest in metrics such as “tokens per watt” and “tokens per infrastructure dollar.”

The appeal is straightforward.

AI data centres exist to generate intelligence, inference and business outcomes. Measuring how much useful AI output is produced for a

given energy or capital investment provides a more direct view of system effectiveness.

Rather than simply minimising energy losses, operators can optimise around the ultimate objective: maximising useful computation within available power constraints.

This perspective encourages more holistic decision-making across compute architecture, cooling technologies, workload scheduling and infrastructure deployment.

For power electronics engineers, such metrics could significantly influence future system design priorities. Power conversion efficiency will remain critical, but evaluation criteria may increasingly incorporate the value generated by every watt consumed.

Accelerated Computing and the Future of EDA


The connection between EDA and accelerated computing continues to deepen.

However, this trend did not begin with AI.

The industry has already navigated previous transitions from single-core processors to multi-core architectures and distributed computing environments. Each transition required new algorithms and software methodologies while delivering corresponding gains in productivity.

Today's shift toward GPUs and specialised accelerators represents the next stage of that evolution.

According to Cadence, accelerated computing has become inseparable from leading-edge EDA development. Platforms such as NVIDIA's CUDA provide abstraction layers that enable software developers to exploit rapidly evolving hardware architectures without continuously rewriting application code.



This significantly accelerates innovation while improving performance across a broad range of engineering workloads. Importantly, GPUs represent only one component of a much larger computational ecosystem.

Purpose-built silicon, custom accelerators and specialised hardware architectures will all play roles in supporting future engineering workflows. The challenge lies in matching the right workload to the right computational resource.

As agentic AI becomes more capable, such optimisation may increasingly be performed automatically.

Closing the Sim-to-Real Gap

One of the most compelling areas for future AI development lies beyond semiconductors themselves. Physical AI, including robotics, autonomous vehicles and intelligent machines, depends heavily on accurate simulation.

The challenge is often described as the "sim-to-real" gap: the difference between behaviour observed in simulation and behaviour encountered in the real world. For autonomous vehicles, simulation already enables billions of kilometres worth of synthetic

testing across diverse weather conditions, environments and traffic scenarios.

Robotics introduces an even more demanding challenge.

Unlike vehicles, which are often attempting to avoid collisions, robots must interact physically with objects. Grasping, manipulating and assembling components require precise modelling of contact forces, material properties and environmental interactions. This so-called "last millimetre problem" demands substantially higher simulation accuracy.

High-fidelity multi-physics environments are therefore becoming essential for generating the synthetic datasets required to train future AI systems. Without them, scaling robotics beyond limited applications may prove impossible.

For EDA vendors, this creates a natural extension of existing expertise. After all, simulating and optimising complex systems has always been a core competency.

Remaining Challenges

Despite the excitement surrounding agentic AI, significant challenges remain.

Unlike consumer AI applications, semiconductor design involves relatively small communities of highly specialised experts. There are not billions of publicly available data points. Instead, critical design knowledge is often proprietary, closely guarded and distributed across numerous organisations.

Photonics, quantum computing, advanced packaging and heterogeneous systems all present rich opportunities for future development. Each domain introduces unique modelling requirements, physical phenomena and optimisation challenges.

Perhaps most importantly, data quality remains a central concern.

AI systems are fundamentally constrained by the datasets on which they are trained. Poor-quality data inevitably limits performance. The semiconductor industry faces a unique challenge in this regard. Unlike consumer AI applications, semiconductor design involves relatively small communities of highly specialised experts. There are not billions of publicly available data points. Instead, critical design knowledge is often proprietary, closely guarded and distributed across numerous organisations.

The volume of publicly accessible semiconductor design data is therefore limited compared with many AI

training domains. This raises important questions about how the industry will develop sufficiently rich, high-quality datasets while maintaining IP protection and competitive differentiation. According to Cadence, solving this challenge may become one of the defining stories of AI-driven semiconductor development over the coming decade.

The Road Ahead

Will agentic AI and digital twins become universal across the semiconductor industry?

Probably not.

Engineering has never adopted any single methodology universally. Cost constraints, application requirements and organisational preferences will continue to shape technology choices. Yet the momentum behind these technologies is difficult to ignore. Early adopters are already reporting significant productivity gains, improved quality of results and accelerated

development cycles. As competitive advantages become evident, broader industry adoption becomes increasingly likely.

The semiconductor sector has always thrived by expanding the boundaries of what is possible. Agentic AI appears poised to become the next major tool in that pursuit, transforming EDA from a collection of automated utilities into an intelligent engineering partner.


For power electronics professionals, the implications extend far beyond chip design. From data centre power architectures and advanced cooling systems to robotics, autonomous transportation and physical AI, these technologies are reshaping how complex systems are optimised and deployed.

The industry's next breakthrough may not come solely from better silicon. It may come from creating engineering systems capable of helping engineers imagine what comes next.

PEI POWER ELECTRONICS INTERNATIONAL

BOOK YOUR REPRINT TODAY!

A reprint from **Power Electronics International** amplifies your editorial exposure and provides authoritative third-party validation. It's a powerful tool for **sales and marketing**, helping you showcase innovation to **engineers, designers, and decision-makers** with independent credibility. Use it across **events, pitches, newsletters, and social channels** to reinforce trust, extend visibility, and strengthen your position in the power electronics market.



Contact: Shehzad Munshi
shehzad.munshi@angelbc.com



Gallium oxide MOSFETs:

Why interfaces matter

Developers of gallium oxide power devices should draw on lessons learnt from the pioneers of compound semiconductor MOSFETs

BY IAIN THAYNE FROM THE UNIVERSITY OF GLASGOW

THANKS to its ultra-wide bandgap (UWBG), alongside availability of native substrates, β -Ga₂O₃ is steadily transitioning from an intriguing laboratory material to a serious option for next-generation power electronics. For researchers and technologists pursuing that goal and working across materials, devices, and integration, the central question is no longer whether gallium oxide channel transistors are scientifically interesting, but how – and how quickly – they can be translated into a robust, scalable device platform that aligns with realistic application and qualification requirements.

It is evident that today's gallium oxide MOSFETs remain at an early stage of technical maturity. Threshold voltages tend to lie on the wrong side of zero, subthreshold swings are well above the thermal limit, and there's a paucity of published reliability data. However,

none of these observations are fundamental showstoppers. Rather, they reflect an emerging MOS technology, where progress in channel material development outpaces refinements in interfaces and gate stacks.


Many of the challenges now confronting gallium oxide are neither unprecedented nor unsolved. As is often the case, delving into the journals of yesteryear can speed development.

For today's developers of Ga₂O₃ MOSFETs, valuable insight can be gained from revisiting earlier work on GaAs siblings, where sustained effort was devoted to interface control, gate-stack chemistry, and charge management. Papers by pioneers of GaAs MOSFETs provide a convincing demonstration that high-performance MOS operation is achievable, even

though the industrial demand required for large-scale adoption did not ultimately emerge at the time.

A recurring challenge: the MOS interface

Whenever a semiconductor material is proposed as a MOSFET channel, the focus quickly shifts from bulk properties to the gate interface. This makes much sense. After all, the interface is crucial, as it is where there's a transition from a crystalline semiconductor to an amorphous dielectric, a change that almost inevitably has the potential to introduce electronic defect states that degrade electrostatics, increase variability, and complicate reliability. One need look no further than the historical development of the silicon MOSFET to see that the ability to control this interface will often determine whether a material system advances from demonstration to deployment.



Gallium oxide is at this inflection point. Many reported β -Ga₂O₃ MOSFETs exhibit large subthreshold swings and threshold voltage instability, behaviour consistent with significant charge trapping, occurring at or near the gate dielectric interface. While there have been impressive advances in crystal growth, substrate availability, and doping control, the same cannot be said for interface engineering. Here, progress is slower, mirroring events in earlier compound semiconductor development cycles.


Lessons from GaAs MOSFET research

During the late 1990s and early 2000s, GaAs MOSFETs were intensively investigated for RF and logic applications. The motivation behind this is that GaAs offers outstanding carrier transport. However, this binary also suffers from well-known surface chemistry challenges, historically limiting MOS device operation. Through the focused research of that time, it was shown that these challenges are not fundamental.

One of the milestones of that era was the development of chemically compatible amorphous gate stacks. Introducing ultrathin gallium-oxide-based interlayers and carefully selecting ternary gadolinium-based high- κ dielectrics with related chemistry enabled the fabrication of GaAs MOSFETs with a low interface state density, near-ideal subthreshold swing, and positive threshold voltage. Crucially, these improvements were maintained through realistic fabrication sequences, demonstrating that these interface solutions were not merely laboratory artifacts.

Ultimately, GaAs MOSFETs with these highly engineered gate stacks did not enter large-scale production. But that's not because they failed to meet performance targets – it's because alternative device technologies were a better match to system-level needs, and the cost structures of the time. Nevertheless, the technical lessons learned – particularly regarding interface control – remain directly relevant to today's UWBG devices.

For those currently working on gallium oxide devices, a noteworthy finding from previous GaAs MOSFET development is the potential of appropriately selected ternary gadolinium-based high- κ materials, deposited optimally on Ga₂O₃. More



➤ Figure 1. Standard gate dielectric screening methods probe defects/traps with energies across virtually all of the bandgap in silicon-based devices (left). In contrast, for Ga₂O₃-based devices (right), trap/defect states with a wide range of energies around mid-gap go undetected. As a consequence, dielectric trap state densities will be underestimated. This is why transistor sub-threshold swing is a better metric for dielectric screening in Ga₂O₃ MOSFETs.

generally, though, there's an underlying philosophy to take onboard – identifying amorphous dielectrics, engineered specifically for chemical compatibility with Ga₂O₃, deserves attention. Unfortunately, much current work involves incremental surface treatments, followed by deposition of conventional dielectrics. While this strategy delivers measurable improvement, it may fail to provide the degree of interface control needed for stable enhancement-mode operation, low variability, and acceptable lifetime under bias.

Switching to more fundamentally engineered gate stacks, combined with appropriate interlayers, promises to suppress interface trap formation systematically, rather than reducing their impact indirectly. For UWBG semiconductors, where deep traps exhibit slow dynamics and long recovery times, this distinction is particularly important from a device-physics and qualification perspective.

Subthreshold swing: A practical indicator

In gallium oxide MOSFETs, subthreshold swing is a particularly useful metric for assessing interface quality. In both wide and ultra-wide bandgap materials, many electrically active states lie deep in the bandgap, making them difficult to probe with the conventional capacitance-based techniques, which were developed to screen oxides for silicon MOS devices as illustrated in Figure 1. Subthreshold swing, by contrast, reflects the cumulative impact of these states on channel formation and electrostatic control.

Based on this state of affairs, as gallium oxide devices mature, convergence toward lower and more reproducible subthreshold swing values will be a key indicator that interface-related challenges are truly being addressed. From an application standpoint, improvements in this metric are closely linked to threshold voltage stability,

Switching to more fundamentally engineered gate stacks, combined with appropriate interlayers, promises to suppress interface trap formation systematically, rather than reducing their impact indirectly

device-to-device uniformity, and ultimately manufacturability.

Architecture considerations

Beyond gate materials, device architectures play an important role. During the development of GaAs MOSFETs, heterostructure channels were used to spatially separate carriers from the oxide interface. This led to reduced carrier interaction with interface states. Similar ‘thin spacer layer’ concepts are now being explored in gallium oxide, through the use of aluminium- and indium-containing oxide alloys.

While one should not underestimate the materials challenges associated with such heterostructures, if successful implementation is realised, this could complement advances in gate stack design. In turn, this could ease interface requirements and expand the available design space for gallium oxide devices.

Towards application readiness

The path from promising research devices to application-ready power transistors is defined by more than

headline performance metrics. As technologies move closer to deployment, it becomes increasingly important to consider the likes of variability, bias stability, endurance under stress, and reproducibility across a wafer. Tightly coupled to each of these considerations is interface quality.

In the near term, if there is meaningful market penetration for gallium oxide MOSFETs, this will probably be associated with grid-level power infrastructure, solid-state protection, pulsed-power systems, and specialised industrial or aerospace platforms. For these applications, scaling volume is not a standout priority. What also matters is realising sufficient confidence in gatestack reliability, having thermal management strategies, and achieving long-term stability under realistic

operating conditions. Success in these early application spaces would represent an important step toward application readiness, establishing gallium oxide as a credible UWBG option for demanding voltage regimes, and laying the groundwork for broader adoption as system requirements evolve.

Development of the GaAs MOSFET shows that it's possible to produce high-quality MOS interfaces on compound semiconductors by addressing chemistry, structure, and processing holistically. Optimisation of the gallium oxide MOSFET may benefit from leaning into these approaches, and leveraging knowledge surrounding gadolinium-based high-κ dielectrics deposited on Ga₂O₃, as well as the judicious use of hetero-interfaces.

FURTHER READING

- M. Passlack *et al.* “Gallium oxide in field effect transistors – a perspective,” *Journal of Vacuum Science & Technology A* **44** 038503 (2026)

PEI POWER ELECTRONICS INTERNATIONAL

BOOK YOUR REPRINT TODAY!

A reprint from **Power Electronics International** amplifies your editorial exposure and provides authoritative third-party validation. It's a powerful tool for **sales and marketing**, helping you showcase innovation to **engineers, designers, and decision-makers** with independent credibility. Use it across **events, pitches, newsletters, and social channels** to reinforce trust, extend visibility, and strengthen your position in the power electronics market.




Contact: Shehzad Munshi
shehzad.munshi@angelbc.com



WÜRTH ELEKTRONIK MORE THAN YOU EXPECT

ULTRA LOW LOSSES

WE-MXGI



© ei50s



With the WE-MXGI Würth Elektronik offers the newest molded power inductor series. It combines an innovative iron alloy material that provides high permeability for lowest R_{DC} values combined with an optimized wire geometry.

Ready to Design-In? Take advantage of personal technical support and free samples ex-stock.

www.we-online.com/WE-MXGI

Highlights

- Extremely high power density
- Ultra low R_{DC} values and AC losses
- Magnetically shielded
- Optimized for high switching frequencies beyond 1 MHz

#UltraLowLosses

Innovations in dynamic switch power testing

How Cosmic is redefining reliability testing for wide-bandgap semiconductor devices at production scale.

BY ROBERT PULMAN, PRODUCT MANAGER - HIGH POWER TEST COSMIC GROUP

THE ELECTRIFICATION of mobility and the rapid proliferation of industrial power conversion is placing unprecedented demands on power semiconductor devices. Silicon Carbide (SiC) MOSFETs and Gallium Nitride (GaN) transistors have emerged as the materials of choice for next-generation power electronics thanks to their ability to switch at higher frequencies, block higher voltages and operate at elevated temperatures compared to conventional silicon. But these performance advantages come with a critical manufacturing challenge: how to verify devices as truly reliable before they reach their end applications? The answer lies in dynamic switch power testing — a discipline that has evolved dramatically in recent years to keep pace with the technical complexity of wide-bandgap (WBG) devices and the tightening requirements of automotive and industrial qualification standards. Cosmic Group, a global partner in advanced semiconductor test solutions, has been at the forefront of this evolution, developing a portfolio of test platforms that bring production-grade dynamic testing to every stage of the device lifecycle from bare dies to packaged module.

Why power device testing cannot be static

Traditional parametric testing — measuring static electrical parameters such as on-resistance, threshold voltage, and leakage current — remains essential, but it captures only a snapshot of a device's behaviour under controlled, steady-state conditions. In real-world applications, power devices are subjected to continuous high-frequency switching, large transient currents, and extreme thermal cycling. Static tests alone cannot reveal how a

device will behave under these conditions, nor can they expose latent defects that only manifest under dynamic stress.

There are four primary motivations for dynamic power device testing, each addressing a distinct aspect of device integrity:

- **Doping and crystal quality** — Verifying wafer quality at the material level to ensure defect-free crystal structure
- **Stress testing** — Driving devices harder than real-world conditions to expose early-life failures
- **Gate oxide layer quality** — Assessing the integrity of the gate dielectric, which is especially critical in SiC devices
- **Die sorting** — Categorising Known Good Die (KGD) parameters to ensure matched performance across dies in multi-die modules

The third and fourth points deserve particular attention for SiC technology. SiC MOSFETs rely on an extremely thin gate oxide layer — the SiO_2 dielectric interface between the Gate and the conduction channel — that is inherently more vulnerable to time-dependent dielectric breakdown (TDDB) than equivalent silicon devices. Carbon migration from the SiC substrate into the oxide layer can cause charge trapping, which shifts threshold voltages and degrades switching performance over time. As voltage ratings climb beyond 1.7 kV in emerging applications, and as switching speeds increase with the adoption of low-inductance packaging techniques, the risks associated with gate oxide integrity become even more pronounced. Dynamic testing is the only way to stress these mechanisms under conditions that replicate actual operation.

The regulatory imperative: AQG324 and evolving automotive standards

The Automotive industry has long operated rigorous qualification frameworks, and industrial power electronics are no exception. The AQG324 standard defines a comprehensive suite of qualification tests, including power cycling HTS, LTS, HTRB, HTGB and H3TRB.

► Cosmic DS5-400 Adapter



Critically, AQG324 mandates dynamic testing as part of the reliability assessment process. Two tests in particular — Dynamic Reverse Bias (DRB) and Dynamic Gate Stress (DGS) — are required for technology qualification, with a minimum of 10^{11} cycles. While this cycle count may seem excessive for production testing, truncated versions of these tests, typically 10k to 100k cycles, are being increasingly requested at final test. At a switching frequency of 1 MHz, a 10k-cycle DRB test completes in just 10 ms — a negligible impact on overall test throughput, especially when index-parallel test configurations are employed. This evolution in standards reflects a broader industry consensus: that WBG devices intended for safety-critical automotive applications cannot be adequately qualified through static testing alone.

Short-Circuit testing: surviving the worst case


One of the most demanding tests in the dynamic testing toolkit is the Short-Circuit Test (SCT). In synchronous motor drive applications, power devices are invariably arranged as three half-bridge configurations. If both upper and lower devices in a bridge leg turn on simultaneously — called “shoot-through” — the result is a direct short-circuit across the battery or DC bus. The power devices must survive long enough for the motor drive protection circuitry to detect the fault, disconnect the energy source and prevent catastrophic damage to the system.

During a short-circuit event the current through the device is limited only by its RDS(on) resistance and the stray inductance of the circuit. For a high-current SiC device, this can mean short-circuit currents reaching up to 10x the normal operating current — a brutal stress test by any measure. SiC devices also exhibit fundamentally different failure mechanisms compared to silicon under high-stress conditions, which makes the SCT an essential part of building a complete picture of device behaviour in production, not just in qualification laboratories.

The critical role of inductance

Implementing SCT in production is not trivial. The test requires a high-energy DC source, fast gate drive circuitry for the device under test (DUT), accurate current and voltage measurement and robust fault-condition protection. But one factor above all others determines the quality and repeatability of the test: stray inductance in the test circuit.

By way of example, measurements on a SiC TO247-4 device with a datasheet short-circuit current of 1,800A at 470V illustrate the impact dramatically: With a 20cm cable connection (~200nH) and a Vds of 100V, the peak current only reaches about 1,000A and the overshoot voltage peaks at 470V. A hard-dock connection (~40nH) on the other hand achieves a peak current of 1,500A and peak Vds of a much more controlled 245V (see adjacent graph.)



The conclusion is unambiguous: reducing inductance in the test path is not a refinement but fundamental requirement for obtaining meaningful SCT results. A hard-dock connection to the DUT, minimising the length of the switching loop, is the preferred approach in production test systems designed for this purpose.

➤ Cosmic 0cm waveform.

Dynamic Reverse Bias testing: Stressing the gate oxide

Where SCT probes the device's ability to survive extreme overcurrent conditions, the Dynamic Reverse Bias (DRB) test targets a different vulnerability: the integrity of the device under repeated high-voltage transient stress. The DRB test applies repeated high-voltage spikes to the drain of the DUT while it is held in the off-state, creating fast charge and discharge cycles through the internal structure of the die. This cycling is specifically designed to expose premature aging mechanisms including charge trapping at the SiC/SiO₂ interface, dielectric breakdown in the gate oxide and degradation of the body diode. AQG324 specifies a minimum slew rate of 50 V/ns for the drain voltage transient, but in practice, with careful attention to reducing the capacitance of the circuit path between the tester and the device handler, slew rates of 100 V/ns and above are achievable and preferable — offering accelerated exposure of latent defects.

A practical DRB test example using a high voltage SiC MOSFET demonstrates the key parameters in action: a switching frequency of 500 kHz and dv/dt of 96V/ns is achievable in production with optimised gate drive conditions. Factors govern the achievable slew rate: the choice of upper switching device and overall capacitance of the drain-source loop, the switching speed of the upper device gate driver, loop inductance, stability of the DUT gate-source voltage and bandwidth of the measurement system.

KGD Testing: protecting module economics

As SiC power modules for automotive applications become increasingly sophisticated — integrating multiple dies in half-bridge and full-bridge

Platform	DS5-400	DS5-QUASAR	DS6-PULSAR
Target Application	Entry-level GaN device testing including dynamic RDS(on)	High-power Si/SiC device testing, die and packaged	Ultra-high-power SiC automotive testing, die and packaged
Switching Current	400 A	1,200 A	3,000 A
Short-Circuit current	500A	2,000A	10,000A
DC Link Voltage	1 kV	1.2 kV	1.4 kV
Parasitic Inductance	Typically 15 nH	Typically 20 nH	Typically 20 nH

configurations — the economics of die-level testing have become compelling. A typical mid-volume automotive-grade 1,200 V SiC half-bridge module has a cost structure in which the individual dies account for approximately 55% of total cost, the substrate for 15%, and final module assembly for the remaining 30%.

Even for the simplest full-bridge module incorporating just 6 dies, the arithmetic is stark: if just one die is defective and the fault is not detected until after module assembly, 93% of total module cost has already been committed. This makes KGD testing a critical process step, not merely a quality assurance formality. Beyond defect screening, parametric KGD testing also enables die-to-die matching, ensuring that all elements of a finished module switch in harmony and share the load symmetrically.

To achieve a module yield of at least 85%, the required average die yield is 99.55%, assuming a defect-free assembly process — a target that demands extremely sensitive and reliable die-level test methods.

➤ Cosmic DS6



Safety in dynamic testing: protecting the test socket

High-power dynamic testing at die level introduces further challenge that is often underappreciated: what happens when device fails during the test itself? Dynamic switch test involves high currents and voltages; if the DUT fails destructively, the energy dissipated in the device and through the test fixture can be sufficient to damage not only the socket but also the mechanical handling components.

Effective over-current protection must accomplish three things simultaneously: distinguish a fault event from a normal test transient, limit the energy delivered to the DUT during reaction window and disconnect energy sources as rapidly as possible. In a well-designed system, shutdown can be achieved in under 300 ns from fault detection — a response time that makes difference between a recoverable failure and catastrophe. Cosmic's SocketSafe™ technology, integrated into its higher-end platforms, provides this level of protection for both KGD and packaged device test sockets, reducing cost and downtime.

Cosmic Group's Dynamic Switch test portfolio

Cosmic's Dynamic Switch product family addresses the full spectrum of test requirements from entry-level GaN device characterisation to ultra-high-power SiC automotive module testing. The portfolio currently comprises three platforms:

The DS5-400 is designed for packaged devices and supports single-, double- and multi-pulse test modes, making it well suited for GaN characterisation where high switching speed is the primary concern. DS5-QUASAR and DS6-PULSAR extend the capability to die-level testing and incorporate SocketSafe™ over-current protection, with PULSAR supporting short-circuit test currents up to 10,000 A for the most demanding automotive SiC qualification scenarios. All 3 platforms are designed around an H-bridge configuration that

enables SCT, DRB and inductive load tests within a single flexible test platform.

Crucially, the same test platforms used for laboratory research and characterisation scale seamlessly to high-volume production. This continuity — from first characterisation through to volume manufacturing — ensures that test conditions applied during device qualification are faithfully replicated at production speed, eliminating the risk of systematic differences between qualification data and real-world production performance.

A new standard for power device reliability

The message from the semiconductor test community is clear: as WBG power devices move from niche applications into mass-market automotive and industrial power conversion at scale, the test strategies developed for conventional silicon are no longer sufficient. Dynamic switch testing — encompassing short-circuit testing, dynamic reverse bias testing, and KGD evaluation — is not a premium option reserved for research labs. It is a production necessity.

Cosmic Group's investment in this area, backed by the combined expertise of Test Inspire, Gedec, RoodMicrotec, ipTEST, Microtest and Focused Test,



represents our commitment to the industry's most pressing reliability challenges.

➤ Cosmic Quasar

With over 25 years of accumulated test expertise and a global footprint spanning Europe, the US, and Asia, the group is uniquely positioned to support device manufacturers as they navigate the transition to the next generation of power semiconductors — and to the test infrastructure that will underpin their success.

 ANGELTECH

CS INTERNATIONAL
CONFERENCE

PIC INTERNATIONAL
CONFERENCE

PE INTERNATIONAL
CONFERENCE

AP INTERNATIONAL
CONFERENCE

AIS ANGELTECH INNOVATE
SUMMIT

SAVE THE DATE

12-14 APRIL 2027

Sheraton Brussels Airport Hotel

To find out more about our sponsor and speaker opportunities,
contact us today on: +44 (0) 2476 718970

or email:

info@angeltech.net

Redefining motor drives with 1 MHz GaN switching

James Cannings, CEO at QPT, discusses the world's first 1MHz GaN motor drive, and the recent opening of customer demonstrations of its MicroDyna test platform, now updated with full Field Oriented Control (FOC) and real-time dynamic cogging correction. James also discusses the launch of qDesign, an AI-driven generative design service that programmatically generates, simulates and iterates on QPTs patented qAttach thermal interface layer for any power module, replacing weeks of human-in-the-loop CAD and simulation cycles with AI-generated topologies and automated optimisation.

FOR MORE THAN two decades, motor drive technology has evolved through incremental improvements rather than fundamental architectural change. Silicon devices became more efficient, silicon carbide extended operating voltages and temperatures, and software delivered increasingly sophisticated field-oriented control (FOC). Yet despite these advances, one characteristic has remained remarkably constant: switching frequencies have largely remained around 10 kHz.

According to James Cannings, CEO at QPT, that ceiling is no longer a technological necessity - it is a historical limitation that has constrained the entire motor drive ecosystem.

QPT's latest demonstration platform, MicroDyna, claims to be the first commercial platform capable of running field-oriented control natively at 1 MHz hard-switching. Achieving this required far more than replacing silicon with gallium nitride (GaN). It demanded an entirely new approach to semiconductor packaging, thermal management, gate driving and system architecture.

If the technology delivers on its promise, the implications extend well beyond faster switching. The platform could fundamentally alter motor design, reduce system bill-of-material costs, eliminate expensive shielding, enable precision motion

using commodity motors and provide unprecedented diagnostic capabilities through software-defined sensing.

Speaking with Power Electronics, Cannings outlined why QPT believes the industry is approaching its biggest shift since vector-controlled drives became mainstream.

Breaking the 10 kHz barrier

Traditional motor drives operate between roughly 4 kHz and 16 kHz, with approximately 10 kHz representing the practical operating point for most commercial systems.

While silicon carbide has extended performance, it has not fundamentally changed this operating paradigm.

"You can push silicon carbide higher," explains Cannings. "Perhaps 50 kHz, maybe even approaching 100 kHz in specialist applications, but the motor itself increasingly becomes the limiting factor."

The reason lies in pulse-width modulation. Conventional inverters present the motor with fast voltage transitions that approximate a sine wave. As switching frequency increases, electromagnetic interference, insulation stress, bearing currents and cable losses become increasingly problematic.

Rather than attempting to optimise this architecture, QPT has changed it entirely. Instead of delivering PWM directly to the motor, MicroDyna



➤ MicroDyna running Field Oriented Control natively at 1MHz hard-switching with dynamic cogging correction active to eliminate torque ripple in low-cost motor.

produces a fully filtered sinusoidal voltage output.

“The motor never sees the switching waveform,” Cannings explains. “It simply sees the clean sine wave it actually wants.”

The distinction is significant. Although sine-wave filters have existed for years, their size has prevented widespread adoption. At conventional switching frequencies they can weigh tens of kilograms while introducing additional losses and cost.

Frequency changes everything. Increasing switching frequency by roughly two orders of magnitude allows the passive filter to shrink dramatically.

“A filter that might weigh 25 kilograms at 10 kHz becomes approximately 250 grams at 1 MHz,” says Cannings.

Instead of becoming an expensive specialist accessory, the output filter becomes practical for mainstream drives. The result is an inverter that presents the motor with an almost ideal voltage waveform while simultaneously reducing acoustic noise, electromagnetic emissions and insulation stress.

Why 1 MHz requires more than GaN

At first glance, GaN appears to make 1 MHz switching inevitable. Its intrinsic switching speed significantly exceeds silicon and silicon carbide, making megahertz operation theoretically possible. The practical reality has proved rather different.

“The semiconductor has never really been the limitation,” Cannings says. He references comments made by GaN pioneer Geoff Haynes, co-founder of GaN Systems, who remarked during the company’s early product launches that creating the transistor had been “the easy part”.

“The difficult question,” Cannings says, “was always how you actually use these devices.”

Packaging quickly became the bottleneck. Conventional power modules were designed around silicon. Simply replacing the die with GaN does not unlock its full switching capability.



➤ MicroDyno’s analysis view showing high-resolution data captured without external sensors. Uniquely possible with high SNR from the pure sine wave motor drive.

Cannings compares the situation with a recent analogy used within the semiconductor industry.

“As Infineon recently put it, putting a Ferrari engine into a tractor still leaves you with a tractor.”

The challenge therefore extends across multiple engineering disciplines.

Efficient 1 MHz hard switching demands switching transitions of approximately one to two nanoseconds while managing voltage slew rates exceeding 550 V/ns. That creates simultaneous RF, thermal and packaging problems.

According to Cannings, solving those problems required years of development across custom gate drivers, electromagnetic design and entirely new packaging technologies. “Our half-bridge modules probably have more in common with RF systems than traditional power electronics.”

qAttach: Solving the thermal challenge

Central to QPT’s approach is a proprietary die-attach technology known as qAttach.

Thermal interfaces are often overlooked outside semiconductor packaging circles, yet they frequently dictate long-term reliability and power density. Conventional solder layers introduce significant thermal resistance. Silver

sintering improves thermal performance but requires substantial pressure during manufacturing while still imposing practical thickness limitations.

QPT instead employs an ultra-thin solder layer - potentially below one micron - combined with a complex surface topology that prevents delamination.

The concept is perhaps best illustrated through an unexpected analogy. One customer compared the design to anti-rip materials used in paragliders. Modern paraglider fabrics divide the surface into numerous small sections so that any tear cannot propagate across the wing. Cannings believes qAttach behaves similarly.

Its structured interface prevents cracks from spreading while maintaining an exceptionally thin thermal path. The outcome is thermal performance reportedly up to fifteen times better than conventional die attachment while avoiding the manufacturing pressures associated with sintering.

Beyond motor drives, the approach has immediate relevance to AI data centre power supplies, where thermal density increasingly limits system performance.

Seeing without external sensors

While 1 MHz switching provides cleaner motor drive signals, it also unlocks another capability: vastly improved

observability. Because the motor receives a pure sinusoidal voltage, returning electrical signals are no longer buried beneath switching noise. This enables extremely accurate measurement of back electromotive force (back EMF).

Internally, QPT samples current and voltage at approximately 4 MHz. The resulting signal quality allows the inverter itself to become an intelligent sensing platform. Cannings prefers to describe the capability as “sensing without external sensors.” The distinction matters.

The drive continues to measure electrical quantities internally, but no longer requires large numbers of external sensors distributed throughout the machine. Instead, subtle disturbances become visible through the electrical behaviour already present within the motor. “The quality of the data changes completely,” he explains. That richer dataset becomes the foundation for predictive diagnostics, adaptive control and AI-based analysis.

Eliminating cogging without lookup tables

One of the first public demonstrations of this sensing capability focuses on torque cogging. Cogging torque results from magnetic imperfections within electric motors and normally requires expensive motors, high-resolution encoders or carefully characterised lookup tables to compensate. QPT instead demonstrated dynamic correction using a deliberately inexpensive motor.

By observing disturbances within the back EMF signal, the controller continuously estimates torque ripple and applies corrections in real time. No calibration tables are required. No lengthy commissioning process is needed. The result is precision typically associated with significantly more expensive servo systems. Cannings believes this could dramatically reduce robotics costs.

Rather than specifying premium servo motors costing several hundred pounds together with high-resolution position encoders, manufacturers could potentially achieve similar performance using commodity motors controlled by intelligent electronics. “We’re demonstrating precision motion using a motor that cost around fifty dollars.”

Reducing system costs

The implications extend beyond motor selection. One of the less visible costs in industrial automation is cabling. Fast PWM waveforms require expensive screened motor cables to limit electromagnetic interference.

In large industrial installations, these cables frequently cost more than the inverter itself. Because QPT outputs a clean sine wave, screened cabling may no longer be necessary. For large factories containing hundreds of metres of cable per motor, the savings become substantial.

Robotics offers another compelling opportunity. Each joint within a collaborative robot typically incorporates premium servos, precision encoders and associated electronics.

Cannings estimates that QPT’s architecture could potentially remove between £500 and £1,000 per linkage. Across six-axis robots, those savings rapidly become several thousand pounds per machine.

Longer term, similar thinking could influence electric vehicle motor design. Much of today’s motor engineering exists specifically to accommodate PWM excitation. If motors instead receive pure sinusoidal drive voltages, some of those design compromises disappear. Preliminary analysis suggests this could reduce EV motor costs by around 20%, although Cannings acknowledges automotive qualification cycles will inevitably take considerably longer.

Digital twins meet Edge AI

MicroDyno incorporates another increasingly important capability: integrated digital twin technology. Rather than transporting large demonstration systems, engineers can recreate complete motor behaviour digitally. More importantly, digital twins provide an efficient environment for AI training.

QPT first models normal motor operation across varying loads before injecting simulated faults including bearing defects and electrical abnormalities. Large datasets generated from these simulations are then used to train compact edge AI models. Once deployed, the drive

continuously monitors incoming data. Within milliseconds it identifies abnormal behaviour before passing observations to a fault classifier.

Cannings acknowledges that predictive maintenance itself is not new. Many industrial suppliers already offer AI-assisted diagnostics. The challenge has always been deployment cost. Traditional systems rely upon multiple external sensors throughout the machine. QPT instead extracts comparable information directly from the motor drive. Because the sensing already exists inside the inverter, the incremental hardware cost becomes negligible.

“We’re really exploiting the richness of the electrical data already available,” Cannings says.

Accelerating power module design with AI

Beyond motor drives, QPT recently introduced qDesign, an AI-assisted design environment intended to accelerate semiconductor packaging development. The system addresses one particularly labour-intensive challenge: customising qAttach geometries for different power modules. Each module possesses unique thermal hotspots, mechanical stresses and coefficient-of-thermal-expansion mismatches.

Traditionally, engineers construct highly detailed CAD models, perform finite-element simulations, manually review results and then begin another design iteration. Each cycle can consume weeks.

qDesign instead uses large language models to generate parametric geometry programmatically using Python and Build123D. Importantly, the AI does not immediately produce manufacturing data. Every design passes through deterministic validation stages ensuring manufacturability, assembly feasibility and compliance with production constraints before entering simulation. Human engineers remain firmly within the process. “The AI isn’t replacing engineering judgement,” Cannings emphasises. Instead, it performs the repetitive optimisation work that humans struggle to execute efficiently.

Simulation outputs containing millions of stress values become machine-


HITACHI



Škoda Group

HV LinPak

After establishing itself as the new standard in the railway application with the Low Voltage (Viso=6kV) variant, the LinPak is now offering a High Voltage variant (Viso=10.2kV). While delivering the higher isolation voltage, it keeps the key features that made the LV variant a big success: phase leg configuration with very low stray inductance, high current density, separation of DC and AC terminals for ideal busbar and gate drive design. With these technical features, it accommodates Si IGBT as well as SiC MOSFET dies with blocking voltage 3.3kV, 4.5kV and 6.5kV.



hitachienergy.com/semiconductors

readable optimisation inputs. The AI modifies microscopic surface geometries before generating revised models for further evaluation. Rather than weeks per iteration, optimisation cycles can complete within 24 hours.

Just as importantly, AI explores parameter combinations that human engineers would rarely attempt manually. Although QPT has yet to discover radically unexpected geometries, Cannings believes customer programmes over the coming months will increasingly demonstrate genuinely novel solutions.

Beyond motor drives


Although MicroDyna has attracted attention for motor control, Cannings emphasises that QPT's packaging technologies extend well beyond industrial drives.

AI data centres represent another significant opportunity. Increasing rack power densities demand smaller, cooler and more efficient power conversion hardware. Many of the same innovations enabling 1 MHz motor drives - including advanced die attachment and high-frequency magnetic components - translate directly into data centre power supplies. In this sense, motor drives simply represent the company's initial commercial focus rather than the limit of its technology.

A licensing strategy

Unlike many semiconductor startups, QPT does not intend to manufacture complete power modules. Instead, it plans to license its intellectual property to semiconductor manufacturers and established module suppliers. Cannings believes this provides the most practical route to industrial adoption. Large OEMs remain understandably cautious about relying upon a small startup for critical supply.

Embedding QPT technology within existing global manufacturing ecosystems provides a lower-risk adoption pathway. Collaborative development projects begin with existing customer modules before optimising qAttach geometries and validating performance through accelerated simulation enabled by qDesign. Only then does manufacturing begin.



➤ qDesign iterating on programmatically generated qAttach thermal-interface geometries. Fully AI generated attach layer enables rapid iteration and optimisation.

The bigger challenge: Industry adoption

From a purely technical perspective, Cannings believes the case for sine-wave motor drives is compelling. The more difficult challenge may prove organisational rather than engineering.

Today's industrial ecosystem evolved around conventional PWM drives. Motor manufacturers, drive suppliers, cable vendors and systems integrators often operate as separate businesses with different commercial incentives. A technology capable of reducing costs across the entire system does not necessarily benefit every participant equally. Drive manufacturers may have spent decades reducing inverter costs by a few cents. Introducing an architecture that fundamentally changes cable requirements, motor design and system economics demands new commercial thinking.

Automotive companies, Cannings argues, may prove more receptive because they already optimise complete vehicle efficiency rather than individual subsystem performance. Industrial automation could require a longer educational process.

Looking ahead

Every generation of power electronics eventually reaches an inflection point where incremental improvement gives way to architectural change. Silicon carbide extended operating

voltage. GaN extended switching speed. QPT argues that the industry's next transformation comes not simply from faster semiconductors but from redesigning the entire system around what those semiconductors make possible.

If successful, motors may eventually be driven exclusively by clean sinusoidal waveforms rather than approximated PWM signals. Acoustic noise could disappear. Cable costs could fall. Commodity motors could perform like precision servos. Motor drives could become intelligent sensing platforms. And semiconductor packaging - traditionally hidden beneath the surface - could emerge as one of the industry's most important enabling technologies.

Whether this vision becomes mainstream remains to be seen.

Industrial qualification cycles are notoriously conservative, particularly within automotive and mission-critical applications. Yet as power density, efficiency and system intelligence continue to dominate engineering priorities, QPT's work illustrates an increasingly important point.

The next revolution in power electronics may not come from inventing a faster transistor. It may come from finally building the system architecture that allows those transistors to fulfil their potential.



12-14
APRIL
2027

Sheraton Brussels Airport Hotel

To find out more about our sponsor and speaker opportunities, contact us today on:

+44 (0) 2476 718970

info@peinternational.net

To find out more about our sponsor and speaker opportunities, contact us today on:

+44 (0) 2476 718970

info@peinternational.net

peinternational.net



Rethinking AI power delivery

Hans Hasselby-Andersen, CEO of Lotus Microsystems, explains why legacy power architectures can no longer meet the demands of modern AI accelerators. The company's vStrata platform combines electrical, thermal and mechanical optimisation into a single low-profile design, delivering power beneath the processor while managing heat at the source.

THE EXPLOSIVE GROWTH of artificial intelligence has forced every layer of computing infrastructure to evolve at an unprecedented pace. Processor performance continues to advance, memory bandwidth is increasing, and cooling technologies have become increasingly sophisticated. Yet one of the most fundamental challenges facing AI infrastructure lies in a less visible part of the system: power delivery.

As AI accelerators move into the kiloampere current regime, conventional board-level power architectures are approaching their physical limits. Longer electrical paths, increasing resistive losses, rising thermal densities and shrinking board space are combining to create bottlenecks that can no longer

be solved simply by incremental improvements.

According to Lotus Microsystems, the industry has reached a point where traditional approaches to power distribution are fundamentally mismatched to modern AI hardware. Its recently announced vStrata platform proposes a complete architectural rethink, replacing conventional horizontal power delivery with a vertically integrated approach that positions power conversion directly beneath the processor.

Rather than presenting another incremental voltage regulator improvement, the company argues that the future of AI infrastructure depends on treating electrical delivery, transient response and thermal

management as a single engineering problem.

When legacy architecture meets modern AI

For decades, horizontal power delivery has served computing systems well. Voltage regulation has traditionally been positioned away from the processor, with power distributed across multilayer printed circuit boards before reaching the point of load. That model worked because processor current requirements remained within practical limits. Today's AI accelerators represent an entirely different class of device.

Modern XPU's routinely require thousands of amps during peak workloads, creating enormous demands on every stage of the power delivery

network. Transporting these currents across conventional PCB traces inevitably introduces resistance, energy loss and heat generation before power even reaches the processor itself.

From Lotus Microsystems' perspective, this represents an architectural rather than a component-level problem.

Instead of attempting to optimise increasingly inefficient horizontal current paths, vStrata effectively removes them. By placing the power delivery module directly beneath the AI processor, electrical path lengths are reduced to almost zero, eliminating much of the resistive loss associated with conventional board-level routing.

The concept appears deceptively straightforward, yet its implications extend well beyond improving converter efficiency. It fundamentally changes how power is distributed throughout the system.

Breaking down the power-thermal divide

Perhaps the most significant aspect of the vStrata philosophy is that it rejects the traditional separation between power engineering and thermal engineering.

Historically, these disciplines have been treated independently. Power delivery engineers optimise conversion efficiency and electrical performance, while thermal specialists address the resulting heat through heatsinks, airflow or liquid cooling. Integration often occurs relatively late in the design cycle.

That separation becomes increasingly problematic as power densities continue to climb.

At hyperscale AI power levels, electrical and thermal performance are no longer independent variables. Every watt lost during conversion immediately becomes heat that must be removed. Every increase in operating temperature influences efficiency, reliability and long-term system performance.

Rather than viewing thermal management as a downstream problem, vStrata integrates both functions within the same physical structure.

Central to this approach is Lotus Microsystems' silicon power interposer technology. Electrically, it provides the high-performance interconnect required for vertical power delivery. Thermally, silicon's excellent conductivity creates a direct pathway for removing heat from the source before hotspots can develop.

For system designers, this represents a shift from balancing competing electrical and thermal compromises toward working with a single co-engineered platform in which both characteristics are optimised simultaneously.

Beyond converter efficiency

One of the headline claims surrounding the vStrata platform is a reported 50% reduction in power conversion losses.

While such improvements naturally attract attention, Lotus Microsystems attributes the gain to a combination of architectural decisions rather than a single technological breakthrough.

The first contribution comes directly from the vertical topology itself.

Reducing electrical path lengths dramatically decreases resistive losses that traditionally accumulate between the voltage regulator and processor. Less resistance translates directly into higher overall efficiency.

The second contribution comes from redesigning the converter architecture itself.

According to the company, conventional integrated voltage regulators encounter an efficiency ceiling as current levels increase. Rather than refining existing architectures, vStrata employs a fundamentally different converter topology intended to bypass those inherent limitations.

Interestingly, the silicon interposer itself is not presented as the primary source of efficiency improvement. Instead, its role is enabling sustained operation at these efficiency levels by managing the remaining thermal load.

Even if conversion losses are reduced by half, systems delivering kiloampere-class currents still generate significant heat. Efficient thermal extraction therefore remains essential for maintaining stable operating conditions and preventing localised thermal hotspots.

The result is an architecture in which electrical efficiency and thermal transport reinforce one another rather than acting as competing constraints.

Meeting the transient challenge

Steady-state efficiency represents only one aspect of AI power delivery.

Modern AI processors routinely transition between workloads that demand enormous current changes within nanoseconds. Any delay in supplying that current risks voltage droop, performance throttling or, in extreme cases, processor instability.



Traditional designs address this challenge by surrounding processors with large banks of external capacitors that act as temporary energy reservoirs during rapid load transitions. While effective, these capacitor arrays consume valuable board space while introducing additional electrical distance between stored energy and the processor itself.

Lotus Microsystems has taken a different approach

Rather than relying on external capacitance distributed around the processor package, vStrata integrates the required energy reservoir directly within the power delivery module.

Positioning this energy storage immediately beneath the processor significantly reduces response latency while simplifying board layout.

The company reports transient response capability exceeding 10 A/ns, allowing the integrated power system to respond rapidly enough to satisfy the sudden current demands characteristic of large AI accelerators. Beyond the electrical performance benefits, eliminating extensive capacitor banks also releases valuable PCB real estate—an increasingly important consideration as servers continue to integrate more compute resources into fixed rack dimensions.

Preventing heat rather than fighting it

Thermal management has become one of the defining challenges of hyperscale AI deployment.

Liquid cooling, advanced cold plates and increasingly sophisticated airflow strategies have all emerged to manage escalating processor power levels. Yet these approaches largely share the same philosophy: remove heat after it has already been generated. Lotus Microsystems advocates a different perspective.

Rather than focusing exclusively on cooling technology, the company argues that the

greatest opportunity lies in preventing unnecessary heat generation in the first place.

Lower conversion losses naturally produce less waste heat

The remaining thermal energy is then transported away through the silicon power interposer, which functions as a highly conductive thermal pathway immediately beneath the processor. Instead of allowing hotspots to form before dissipating them through external cooling systems, heat begins moving away from its source almost immediately after it is generated. The company reports point-of-load temperature reductions of up to 25°C. Such improvements extend beyond simply lowering processor temperatures.

Reduced operating temperatures typically improve long-term component reliability, increase expected lifetime and provide greater thermal headroom for future processor generations. Lower baseline temperatures may also enable greater compute density within existing rack envelopes by reducing cooling constraints.

Viewed from a system perspective, the strategy effectively combines two complementary mechanisms: generating less heat during power conversion while simultaneously improving the efficiency with which the remaining heat is removed.

Integration without disruption

Technical superiority alone rarely guarantees adoption within the semiconductor industry.

Infrastructure vendors have

invested heavily in existing reference designs, controller ecosystems and development workflows. Any new architecture requiring wholesale redesign presents a significant barrier to deployment.

Lotus Microsystems appears to have recognised this challenge from the outset.

Although vStrata introduces a fundamentally different physical power architecture, the platform has been designed to remain compatible with existing Tier 1 reference designs and commercially available power management controllers.

Rather than asking processor vendors and system integrators to rebuild their entire design ecosystem, the objective is to integrate within existing development environments.

This compatibility strategy could prove particularly important during early adoption. Silicon vendors can continue using familiar controllers, established design methodologies and proven software infrastructure while evaluating the performance advantages offered by vertical power delivery.

Reducing engineering friction may ultimately prove as important as improving efficiency.

Collaboration before tape-out

Another notable aspect of the programme has been the level of collaboration preceding silicon tape-out.

Rather than developing the architecture independently before approaching customers, Lotus Microsystems reports extensive engagement with leading XPU and AI



At hyperscale AI power levels, electrical and thermal performance are no longer independent variables. Every watt lost during conversion immediately becomes heat that must be removed. Every increase in operating temperature influences efficiency, reliability and long-term system performance.

infrastructure partners throughout the design process.

This collaborative approach appears to have shaped the platform's development around practical deployment challenges rather than theoretical optimisation.

According to the company, early feedback following tape-out has reinforced confidence that the platform addresses genuine bottlenecks facing hyperscale AI designers.

Among the most frequently cited advantages has been improved compute density.

Rack space remains one of the most valuable resources within hyperscale data centres. Any technology capable of reducing the physical footprint of power delivery while simultaneously lowering thermal constraints creates opportunities to deploy additional compute resources within existing infrastructure.

This directly influences total cost of ownership.

Higher compute density improves utilisation of expensive facilities, networking infrastructure and cooling systems while potentially reducing overall capital expenditure for equivalent AI capability.

For hyperscale operators deploying tens of thousands of accelerators, even modest gains at the server level can translate into substantial operational savings.

The road to deployment

Although vStrata represents an ambitious architectural shift, commercial deployment will necessarily proceed in stages.

Engineering samples are scheduled to enter an early access programme during the third quarter of 2026, allowing selected partners to transition

from theoretical evaluation toward practical system integration. This phase will provide the first opportunity to validate laboratory performance under real-world AI workloads and hyperscale operating conditions.

Following engineering sampling, the focus shifts toward collecting operational data from pilot deployments before progressing through qualification and eventually entering volume production.

These milestones will determine whether the architecture delivers its promised benefits outside controlled demonstration environments. For the broader industry, the coming year may prove particularly significant. If vertical power delivery demonstrates measurable improvements in efficiency, thermal performance and compute density within production-scale AI systems, it could influence how future accelerators are architected from the outset.

A new direction for AI power electronics

The demands of AI have repeatedly forced engineers to reconsider long-established assumptions about computing architecture.

Memory hierarchies have evolved, cooling technologies have transformed and interconnect bandwidth continues

to increase. Power delivery may now be approaching a similar inflection point. The challenges facing modern AI infrastructure are increasingly driven by physics rather than incremental optimisation. Transporting kiloampere currents across conventional PCB layouts while maintaining efficiency, thermal stability and transient performance becomes progressively more difficult as processor power continues to rise.

vStrata represents one possible answer to that challenge—not through incremental refinement, but through architectural reorganisation. Its central proposition is that electrical delivery and thermal management can no longer be treated as independent disciplines. Instead, they must become a unified design problem in which converter topology, energy storage, interconnection and heat removal are engineered as a single integrated system.

Whether vertical power delivery ultimately becomes an industry standard remains to be seen. However, as AI infrastructure continues its relentless scaling trajectory, the underlying argument appears increasingly compelling. Future advances in processor performance may depend not only on faster silicon, but on fundamentally rethinking how electrical energy reaches it in the first place.





Rethinking AI data centre cooling

Brooks Henderson of Phononic explains why cooling must be viewed as part of the power delivery and efficiency equation, with thermal management emerging as a core design parameter alongside power and performance.

THE AI REVOLUTION has created a new bottleneck in computing infrastructure, and it is not necessarily processing power.

For decades, semiconductor innovation focused primarily on delivering more transistors, greater bandwidth and higher computational throughput. Today, however, the industry's ability to exploit those advances increasingly depends on something far more fundamental: heat.

As GPUs approach kilowatt-level power consumption, high-bandwidth memory (HBM) stacks become denser, and co-packaged optics move closer to compute devices, thermal management is emerging as one of the defining engineering challenges of the AI era.

According to solid state cooling specialist Phononic, thermal physics is now setting many of the practical performance limits for AI infrastructure. The company argues that the next major gains in compute performance, energy

efficiency and infrastructure utilisation may come not from ever more powerful silicon alone, but from smarter and more precise thermal control.

The march toward the megawatt rack

The scale of AI infrastructure is changing at an extraordinary pace. What only recently were considered high-density racks consuming tens of kilowatts are rapidly giving way to systems measured in hundreds of kilowatts. Looking further ahead, the industry is actively discussing one-megawatt AI racks as next-generation deployments push power densities to unprecedented levels.

That growing power demand ultimately feeds a common goal: delivering more AI performance.

As power is distributed from the rack level down to individual devices, its consumption is concentrated in increasingly demanding GPUs, CPUs, networking silicon and memory

subsystems. The result is a steady increase in thermal density throughout the data centre.

Modern AI accelerators already consume several hundred watts, with future generations moving closer to the kilowatt range. At the same time, memory bandwidth requirements continue to rise, driving ever more sophisticated HBM architectures and advanced packaging technologies.

The problem is straightforward. Semiconductor devices are designed to operate within tightly controlled temperature envelopes. Once those limits are approached, protective mechanisms begin to take effect.

GPUs can reduce operating frequencies. Memory devices can enter protective refresh modes. System performance is throttled to prevent degradation and preserve reliability.

In other words, the silicon itself is often capable of delivering more

performance than the thermal environment allows.

Current cooling strategies largely address this challenge through over-provisioning, employing increasingly sophisticated air and liquid cooling infrastructure to ensure acceptable operating conditions under worst-case scenarios.

Phononic believes that approach is becoming increasingly inefficient.

Targeting the hotspots

The company's alternative strategy focuses on what it describes as intelligent, localised thermal management.

At the heart of its approach is solid state thermoelectric technology, which differs significantly from conventional mechanical cooling systems.

Unlike compressor-based cooling solutions, thermoelectric devices can both cool and heat. This capability may seem counterintuitive in an AI data centre, but maintaining components within their optimal thermal operating window often requires precise temperature control rather than simply achieving the lowest possible temperature.

Just as importantly, thermoelectric devices are themselves semiconductor components.

This creates opportunities for integration and telemetry that are difficult to achieve with traditional bulk cooling systems. Temperature data, operating conditions and system performance information can be captured directly at the point where thermal management occurs.

For data centre operators, this introduces a new level of visibility into how compute, memory and optical subsystems are behaving in real time.

Rather than viewing cooling as an entirely separate infrastructure layer, the technology effectively embeds thermal intelligence directly into the computing environment.

The HBM challenge

Nowhere is the thermal challenge more apparent than in high-bandwidth memory.

HBM has become one of the critical enablers of modern AI systems. Training and inference workloads require enormous amounts of memory bandwidth, driving continuous advances from HBM3 through HBM4 and into future generations.

However, each performance increase brings a corresponding thermal challenge.

According to Phononic, first-generation cooling solutions were designed around heat flux levels typically associated with HBM3 implementations. Future HBM4 platforms are expected to demand significantly higher cooling capability as power densities continue to increase.

The challenge becomes even more pronounced as packaging technologies evolve.

Advanced 2.5D and 3D integration schemes place compute and memory devices in extremely close physical proximity, delivering substantial gains in performance and bandwidth while simultaneously concentrating heat generation.

Traditional system-level cooling methods can struggle to address these localised hotspots efficiently.

Phononic's approach combines thermoelectric cooling hardware with software-defined thermal management. Firmware continuously monitors conditions and dynamically adjusts cooling behaviour according to workload requirements.

This allows operators to optimise for different objectives.

Some applications may prioritise maximum performance. Others may focus on maximising reliability and

component lifespan. Still others may seek the lowest possible energy consumption. The ability to adapt cooling behaviour to specific operational requirements represents an important departure from fixed cooling architectures.

The hidden cost of thermal throttling

One of the most significant consequences of inadequate thermal management is performance loss through throttling.

As HBM temperatures rise, memory devices can transition into protective operating modes designed to prevent long-term damage. While these mechanisms preserve reliability, they also reduce available performance.

For AI workloads, this can have a substantial impact.

Modern training and inference systems are highly dependent on memory bandwidth. Any reduction in memory performance can create bottlenecks that diminish the utilisation of expensive accelerators. Phononic argues that targeted hotspot cooling offers a direct path to improving effective silicon utilisation. By keeping memory devices within optimal temperature ranges, throttling events can be reduced or eliminated. GPUs and memory systems maintain peak performance for longer periods, allowing operators to extract greater value from existing hardware investments.

This relationship highlights an increasingly important shift in data centre thinking.

Raw chip specifications alone no longer define actual performance. Instead, realised performance increasingly depends on maintaining thermal

Modern AI accelerators already consume several hundred watts, with future generations moving closer to the kilowatt range. At the same time, memory bandwidth requirements continue to rise, driving ever more sophisticated HBM architectures and advanced packaging technologies

conditions that allow devices to operate as intended.

From this perspective, cooling becomes an active contributor to computational output rather than simply a background utility.

Beyond device cooling

Improved thermal management also creates opportunities at the system level.

A conventional approach to AI infrastructure often requires highly conservative cooling strategies. Operators maintain aggressively low coolant temperatures to ensure hotspots never approach critical thresholds.

While effective, this approach consumes significant energy.

According to Phononic, targeting localised hotspots can allow operators to increase liquid cooling temperatures elsewhere in the system while maintaining device-level reliability.

The result is the potential for meaningful improvements in Power Usage Effectiveness (PUE) and overall energy efficiency.

Rather than cooling the entire system to accommodate a handful of high-temperature regions, resources can be applied precisely where they are needed most.


For hyperscale operators managing thousands of AI servers, even modest efficiency improvements can translate into substantial operational savings.

Co-packaged optics raise the stakes

Thermal management challenges extend well beyond GPUs and memory.

The rapid emergence of co-packaged optics (CPO) represents another major shift in AI infrastructure design.

As networking bandwidth requirements continue to accelerate toward 1.6T and beyond, traditional pluggable transceivers are increasingly unable to keep pace. The industry is therefore moving optical functions closer to switching silicon through co-packaged architectures.



While this approach addresses bandwidth constraints, it introduces new thermal complexities.

Historically, many optical components resided on switch faceplates where thermal isolation was relatively straightforward. Co-packaged optics bring sensitive photonic devices much closer to high-power electronics.

At the same time, external laser sources are becoming more powerful and generate additional heat loads of their own.

Unlike digital electronics, optical devices often exhibit extreme sensitivity to temperature variation. Small thermal fluctuations can directly affect performance and reliability.

This creates a strong requirement for thermal co-design.

Cooling can no longer be treated as an afterthought once systems have been architected. Instead, thermal considerations must be integrated alongside photonic and electrical design from the outset.

Phononic already has extensive experience in optical temperature control, with millions of thermoelectric cooling devices deployed in laser applications. The company sees this expertise becoming increasingly

relevant as co-packaged optics move into mainstream AI infrastructure.

Rethinking cooling as part of power delivery

For power electronics engineers, perhaps the most important shift is conceptual.

Cooling has traditionally been viewed as an external support system, separate from power conversion and distribution infrastructure.

That boundary is beginning to disappear.

Phononic argues that thermal management should increasingly be considered part of the broader power optimisation equation.

Data centre designers typically operate within fixed power budgets. Every watt consumed by cooling infrastructure is a watt unavailable for computation.

However, intelligent thermal management introduces new possibilities. Faster response times, predictive operation and precise hotspot control allow systems to use cooling resources more efficiently.

The energy saved can either reduce operating costs or be redirected toward additional computational capacity.

This transforms cooling from a passive energy consumer into an active mechanism for power allocation and optimisation.

For power engineers, that means thermal design decisions may soon have direct implications for overall system performance and utilisation.

The rise of the thermal fabric

One of Phononic's most ambitious concepts is what it calls the "Thermal Fabric."

The idea builds on the growing convergence of telemetry, thermal control and data centre orchestration.

Phononic's thermal kits combine solid state cooling devices with onboard firmware capable of monitoring temperature, voltage and operating conditions. These individual systems can then communicate through standard management interfaces and Redfish APIs.

Once connected, thermal management becomes part of a broader infrastructure control system. Instead of optimising cooling device by device, operators gain the ability to optimise thermal performance across entire racks or data centres.

Thermal information can potentially be integrated with workload scheduling, virtualisation platforms and power management systems.

This opens the door to new forms of optimisation.

Workloads could be distributed not only according to compute availability but also according to thermal headroom.

Cooling resources could be allocated dynamically where they provide the highest value. Data centres could make energy decisions based on real-time thermal intelligence.

The result would be a much tighter integration between compute, power and cooling domains than exists today.

Economics drive adoption

For hyperscalers, however, technology must ultimately be justified economically. Phononic believes the strongest case for intelligent cooling

lies in the combination of performance gains and infrastructure efficiency.

Reducing throttling directly increases workload throughput. In AI environments increasingly measured by tokens generated and processed, higher utilisation translates directly into higher economic value.

At the same time, more efficient cooling can reduce operational expenditure and improve facility efficiency.

Importantly, the benefits operate at multiple levels.

Some value is realised at the semiconductor level through improved device utilisation. Some is captured at the package level through hotspot management. Additional gains emerge at the rack and infrastructure level through reduced cooling overhead.

This layered benefit structure is likely to become increasingly important as AI deployments continue to scale.

Thermal management takes centre stage

Looking ahead, Phononic believes thermal management will become a first-class design parameter alongside power and performance.

The trend appears difficult to avoid.

AI systems are simultaneously increasing power density, adopting advanced packaging technologies and integrating sensitive photonic

components. Each development amplifies thermal complexity.

Moreover, hotspots are no longer confined to individual chips. Future infrastructures will contain thermal challenges across GPUs, memory, networking devices, optical components, power supplies and voltage regulation systems.

Managing those hotspots effectively will require faster response times, greater visibility and deeper integration between thermal systems and broader infrastructure management platforms.

In many ways, the industry's traditional view of cooling is already becoming obsolete. The emerging AI data centre is not simply a collection of processors supported by cooling equipment. It is a tightly coupled system in which compute, power and thermal behaviour continuously influence one another.

As the industry advances toward megawatt-scale AI racks and increasingly sophisticated architectures, the ability to optimise those interactions may become one of the most important competitive advantages available.

The next leap in AI performance may therefore come not from adding more silicon, but from ensuring existing silicon can operate at its full potential. In that future, thermal intelligence could prove just as valuable as computational intelligence itself.

[Watch the original interview here.](#)



How to pick a hot-swap MOSFET

A comprehensive guide on how to select MOSFETs for server applications and how to evaluate their suitability for hot-swap operation.

BY NEXPERIA


Hot-swap application overview

HOT-SWAP CIRCUITS are found in applications where a subsystem can be removed without disturbing or powering down the rest of the system. In the case of computing and data management, such as AI servers; individual trays of GPUs, CPUs, or network switches cannot disturb the rest of the server rack within the data center. This makes AI servers an excellent candidate for hot-swap circuits.

As a result, hot-swap circuits require in-rush current control, fault protection, and controlled turn on- and turn-off functions as they pass power from a centralized source (i.e. 48 V busbar) to a subsystem (i.e. individual server trays). The MOSFET pass element (Q1 in Fig. 1 IC (Hot-swap Controller in Fig. 1) and controller) are separated because of high power requirements. As a result, certain MOSFET parameters must be considered.

Maximum drain-to-source voltage (V_{DSmax}), drain-to-source on-state resistance (R_{DSon}), packaging, paralleling technology, and Safe Operating Area (SOA) will be discussed in detail to simplify the MOSFET selection process. In addition, this will explain why Nexperia's enhanced SOA hot-swap MOSFETs, such as the PSMN2R3-100SSE, excel in the market.

➤ Figure 1.
Hot-swap
block diagram.



MOSFET parameters for standard operation

Input voltage (V_{IN_typ}) vs. maximum breakdown voltage (V_{DSmax}) For newer and modern data center architectures, the busbar voltage, or input voltage (V_{IN_typ}) to an individual server tray, is within a range of 40 V – 60 V. The typical value of the input voltage will either be 48 V, 50 V, or 54 V. However, transient voltage surges or spikes, in combination with system parasitics (i.e. busbar and wires from the rack to the power distribution board), will induce voltage much higher than the typical range and destroy the MOSFET if the maximum breakdown voltage (V_{DSmax}) is exceeded.

As a result, most designers have determined that V_{DSmax} is 100 V for 54 V systems. More information can be found in the Calculating V_{DSmax} of a MOSFET section.

Because V_{DSmax} determines packaging, R_{DSon} range, and ultimately price of the MOSFET, it is important to pick the V_{DSmax} first. In any case, Nexperia offers enhanced SOA hot-swap MOSFETs with V_{DSmax} at both 80 V and 100 V, such as the PSMN2R3-100 SSE.

Input power and current vs. MOSFET power dissipation (R_{DSon})


Input power (P_{IN_typ}) is the wattage requirement determined by the circuits within a subsystem. In the case of server trays, GPUs and CPUs will determine most of the power, which ranges from 1.2 kW – 2.2 kW in current architectures and 2 kW – 4 kW in near-future architectures.

Equation 1 Where: determines the typical input current (I_{IN}):

$$I_{IN} = \frac{P_{IN_typ}}{V_{IN_typ}}$$

$$I_{IN} = \frac{(4 \text{ kW})}{(54 \text{ V})}$$

$$I_{IN} = 74 \text{ A}$$



- I_{IN} is the input current seen by the system [A]
- P_{IN_typ} is the typical input power seen by the system [W]
- V_{IN_typ} is the typical input voltage seen by the system [V]

Equation 1 (Equation 1) yields a range of 22 A – 74 A, which are typical input currents in modern server architectures. This ultimately leads to power dissipation (PD) in the MOSFET, which is the limiting factor in most MOSFET circuits.

While the maximum junction (or die) temperature rating of Nexperia hot-swap MOSFETs is 175 °C, most designers will have a target junction temperature (T_{J_target}) and other system requirements that they must adhere to.

Typical target junction temperatures can range from 100 °C – 125 °C to allow for transient events, and ambient temperatures can range from 50 °C – 85 °C depending on cooling paths within the server tray. Thermal resistance (R_{thJA}) is determined by the MOSFET package and layout, which is discussed in more detail in MOSFET Packages and Paralleling with Application Specific (ASFET) technology section, but an 8x8 mm MOSFET with good layout is used for this example.

In combination with the calculated input current (74 A), 0.328 mΩ is calculated in the Calculating target R_{DSon} of a singular MOSFET section. However, there are no 0.328 mΩ MOSFETs available on the market. As a result, designers will place several MOSFETs in parallel to achieve higher current capability.

MOSFET packages and paralleling with ASFET technology

Fig. 2 shows a graph of popular hot-swap

MOSFET's, including PSMN2R3-100SSE, maximum server tray power vs. horizontal PCB area. This determines how many MOSFETs in parallel are needed for a system's target power.

To achieve the target input current of 74 A (or 4 kW), Fig. 2 indicates that several parallel MOSFETs (e.g. 4) can be used. This is determined by Equation 2:

$$R_{DSon_max@temp} = \frac{T_{J_target} - T_{amb}}{\left(\frac{I_{IN}}{n_{parallel}}\right) \cdot R_{thJA}}$$

$$R_{DSon_max@temp} = \frac{(113^{\circ}C) - (50^{\circ}C)}{\left(\frac{(74A)^2}{(4)}\right) \cdot (35^{\circ}C/W)}$$


$$R_{DSon_max@temp} = 5.259 \text{ m}\Omega$$

Where:

- T_{J_target} is the target junction or die temperature [°C] (Equation 2)
- R_{thJA} is the thermal resistance from junction to ambient, which usually consists of PCB and MOSFET package thermal resistance [°C/W or K/W]
- T_{amb} is ambient temperature of the system [°C]
- I_{IN} is the input current seen by the system [A]
- $R_{DSon@temp}$ is the maximum R_{DSon} adjusted from 25 °C to the target temperature using the data sheet [Ω]
- $n_{parallel}$ is the number of MOSFETs in parallel

As discussed in "LFPK MOSFET thermal design guide How to pick a hot-swap MOSFET" application note, package construction and dimensions play a large role in determining thermal impedance or thermal capability of a MOSFET. Both the PSMN2R3-100SSE (LFPK88,

➤ Figure 2. Power Density - Max Power Dissipation vs. Horizontal Length Target ($T_J = 113^{\circ}C$, $T_A = 50^{\circ}C$).



➤ Figure 3. I_D vs V_{GS} for regular and ASFET technology to improve paralleling.


SOT1235, 8x8mm) and PSMN1R0-100ASE (CCPAK, SOT8000A, 12x12mm), benefit from copper clip technology for the best thermal performance, as opposed to bond wires used in TOLL (e.g. PG-HSOF-8-1, 10x12mm). This performance improvement will be reflected in a MOSFET's Safe Operating Area (SOA) graph, and thermal package impedance graph located in the data sheet.

While packaging is important, PCB layout is a main contributor to the thermal path (i.e. 0.44 °C/W for the thermal impedance of package vs. approximately 35 °C/W for the thermal impedance of the PCB). Ideal and realistic PCB thermal resistances (R_{thJA}), such as 35 °C/W for LFPAK88, can be found in the "Thermal boundary condition study on MOSFET packages and PCB substrates" application note.


The PSMN2R3-100SSE's maximum R_{DSon} at 175 °C is 5.2 mΩ, which is a great choice for this application requirement of 5.259 mΩ calculated by Equation 2. This difference adds margin when comparing the actual junction temperature to the target.

Note that Fig. 2 shows that the PSMN1R0-100ASE has the best power density (i.e. largest slope) out of any device in the market. The PSMN1R0-100ASE is a 1 mΩ, 12x12 mm, Enhanced SOA MOSFET using the CCPAK package. This device has a clear advantage and gives the best performance, or most margin, at higher power (i.e. >5 kW designs under the same assumptions).

As discussed in "Paralleling power MOSFETs in high power applications" application note, mismatch in gate-source threshold voltage (V_{GSth}), sometimes called the MOSFET "on" voltage, will lead to uneven



➤ Figure 4. Example of unmatched MOSFET threshold voltages (V_{GSth}) resulting in inequal input current (IIN) for Regular (right) and ASFET (left) technology.



➤ Figure 5. Hot-swap simulation and typical waveforms.

current sharing between MOSFETs that would result in higher input current (I_{IN}) for a singular MOSFET as shown in Fig. 4.

Fig. 3 shows 3 types of MOSFET technology: Regular technology (PSMN2R0-100SSF), Application Specific MOSFET (ASFET) technology with enhanced SOA (PSMN2R3-100SSE), and ASFET technology with enhanced current sharing (PSMN2R3-100SSJ). Choosing enhanced SOA ASFET family will result in 2 to 4 times better current sharing for worst case V_{Gsth} separation compared to Regular technology. This is also demonstrated in the waveforms shown in Fig. 4.

As a result, choosing a MOSFET in the enhanced SOA ASFET family, such as PSMN2R3 100SSE, will benefit from the enhanced current sharing.

MOSFET parameters during hot-swap operation

As seen in Fig. 5, typical hot-swap controllers will linearly and constantly charge the gate of the


MOSFET to facilitate a controlled turn-on. This operation puts the MOSFET into current source mode, sometimes called linear mode, which constantly charges the bulk capacitance (V_{out}). The MOSFET current source mode is extensively discussed in the “Power MOSFETs in linear mode” application note, and one of the most stressful modes a MOSFET can operate in.


As a result, there are 4 important aspects for linear mode operation:

- Voltage at the input to be transferred to the output (V_{IN} or V_D)
- In-rush current needed to charge the bulk capacitance (I_{INRush})
- The amount of time taken to charge the bulk capacitance (t_{pulse})
- Pulse Shapes of current and voltage or energy (triangular or rectangular)

Input voltage (V_{IN}) is the same as the system’s input voltage, so 54 V is used.

➤ Figure 6. Safe Operating Area (SOA) comparison between Regular and ASFET technology.





➤ Figure 7. Example start-up into short circuit waveforms using ADI's LTC4238 hot-swap controllers.

The V_{DS} voltage pulse shape is triangular, and the current pulse is rectangular, which results in a triangular power pulse shape as shown in Fig. 5. This effectively doubles the SOA performance of the MOSFET as discussed in "Power MOSFETs in linear mode" application note.

The amount of time (t_{pulse}) and the in-rush current (I_{INrush}) depend on bulk capacitance, the input voltage (V_{IN}), the hot-swap controller's gate voltage slew rate or source current, and fault conditions. For most hot-swap server applications, this time is between 800 μs – 500 ms. In modern architecture, 10 ms or 100 ms are the typical benchmarks.

Ultimately, these values are compared against the MOSFET's Safe Operating Area (SOA) graph in the data sheet, as shown in the Safe Operating Area (SOA) of a MOSFET and ASFET technology section.

Due to the complex interaction between MOSFET, Hot-swap Controller, and bulk capacitance; simulation is the recommended way to establish a baseline before the prototyping phase. Refer to Fig. 5 and "Nexperia Precision Electrothermal models in SPICE and VHDL-AMS for Power MOSFETs" application note. Otherwise, the Imprecise calculation of hot-swap pulse time will help for generic calculations.

➤ Figure 8. Common hot-swap faults in context of Safe Operating Area (SOA).

Safe Operating Area (SOA) of a MOSFET and ASFET technology

Safe Operating Area (SOA), as shown in Fig. 6, illustrates the range of voltage, current, and pulse length where operating below the lines will indicate safe operation.

For power MOSFETs in linear mode there is a significant reduction in performance, which can be reviewed in the Definition of Safe Operating Area (SOA) regions and Spirito effect section. This phenomenon is known as the Spirito effect, and MOSFETs with Enhanced SOA ASFET technology eliminate the Spirito effect which increases SOA performance and margin as shown in Fig. 6.

Comparing the example in Fig. 6 where V_{IN} is 54 V, the pulse shape is triangular, t_{pulse} is 100 ms, and the MOSFET is the PSMN2R3-100SSE; the SOA current results in 14 A (or 2 multiplied by 7) at $T_J = 25^\circ C$ for, which is approximately 4 times better performance compared to the PSMN2R0-100SSF, non-enhanced SOA MOSFET (3.5 A or 2 multiplied 1.75).


As shown in the Imprecise calculation of hot-swap pulse time section example, the regulated in rush current is around 8 A and 54 V, for 120 mF and approximately 100ms, which is below the 14 A limit at 100 ms for the PSMN2R3-100SSE.

Recalculating SOA performance for different t_{pulses} , shapes, and temperatures is discussed in "Power MOSFETs in linear mode" application note and Example of re-rating SOA graph for temperature section.

Fault conditions such as start up into short circuit

There are many fault conditions that can occur to a MOSFET in a hot-swap application. The most demanding fault is start-up into short circuit as shown in Fig. 7.

Start-up into short circuit shows that the V_{DS} and current are both large which result in a larger



instantaneous power compared to a running short-circuit, where V_{DS} is small but current is large. This is shown on the right side of Fig. 8. In some cases, this energy is more than the energy seen during a normal hot-swap operation to charge the bulk capacitors.

As a result, it is important to compare the V_{IN} , I_{INRush} , t_{pulse} , and shape during this fault condition to the SOA graph in data sheet.

The role of hot-swap controllers during fault conditions

Ultimately, the resulting waveform during a fault condition is the combination of parasitics in the system and the hot-swap controller's reaction. If a MOSFET passes every use-case except a fault, changing the MOSFET might not be the solution or an option without a redesign.

Luckily, the best hot-swap controllers, such as Texas Instrument's LM506x series and Analog Devices LTC42xx series, are configurable using software which can help reduce stress on the MOSFET without redesigning.

For example, the current is detected through the shunt resistor and the sensing pins of the hot-swap

controller as shown in Fig. 1. Fault detection time and deglitch filters of the hot-swap controller will determine when the fault triggers, which should help determine worst case t_{pulse} . Both LM506x and LTC42xx can adjust or shorten these fault timers.

As a result, always consider the role of the hot-swap controller and how it can reduce energy exposed to the MOSFET.

Conclusion

Hot-swap circuits are found in AI server applications because of the relationship between the rack's busbar and the server tray. Hot-swap circuits require features, such as in-rush current control, fault protection, and controlled turn on and turn off that will stress the MOSFET. Key parameters such as maximum drain-to-source voltage (V_{DSmax}), drain-to-source on-state resistance (R_{DSon}), packaging, paralleling technology, and Safe Operating Area (SOA) must be considered in the MOSFET selection process.

Notably, Nexperia's enhanced SOA hot-swap ASFET MOSFETs, such as the PSMN2R3-100SSE, are perfect for modern and near future server hot-swap designs.

PEI POWER ELECTRONICS INTERNATIONAL

BOOK YOUR REPRINT TODAY!

A reprint from Power Electronics International amplifies your editorial exposure and provides authoritative third-party validation. It's a powerful tool for sales and marketing, helping you showcase innovation to engineers, designers, and decision-makers with independent credibility. Use it across events, pitches, newsletters, and social channels to reinforce trust, extend visibility, and strengthen your position in the power electronics market.



Contact: Shehzad Munshi
shehzad.munshi@angelbc.com



From road to rack

A look at how 800V EV innovations are redefining AI data centre power architecture.

BY MATTHEW FALL, TECHNOLOGY ANALYST AT IDTECHEX

THE TRANSITION to 800V in electric vehicles has accelerated the commercialization of wide bandgap semiconductors, and has led to increased vehicle range, faster charging, and greater e-powertrain efficiency. AI data centers are being pushed to ever-increasing power ratings to support future generations of AI training and inference. Drawing inspiration directly from EV power electronics, the data center industry is working to develop the HVDC (800VDC) data center, enabled by wide bandgap power electronics innovations.

The electric vehicle industry has, for the past decade, become the dominant application area for power electronics. Most power electronics players have enjoyed the automotive industry as their main source of revenue for years, and power electronics innovations have focused mainly on improving efficiency and performance across the e-powertrain. However, throughout the past two years, power electronics tier 1s have gradually seen their automotive industry revenue stagnate amid reduced government incentives and increasingly steep competition with highly vertically integrated Chinese EV manufacturers.

While IDTechEx still forecasts long-term growth in the EV market, many players are looking for alternative, growing application areas for power electronics. Data centers, and especially AI data centers, represent a growing and rapidly innovating market, with an increased uptake of wide bandgap semiconductors SiC and GaN to support increasingly powerful and power-hungry AI training models. According to

IDTechEx forecasts in the market report “Power Electronics Market 2026-2036: Data Centers, Electric Vehicles, and Renewables”, the power electronics market for data centers is expected to grow 2.5-fold by 2036.

A paradigm shift in data center power architecture is coming over the next few years in the form of HVDC (800VDC) data centers, which will unlock an order-of-magnitude increase in rack power. This overhaul in data center power architecture is directly influenced by 800V EV power electronics, from the materials used to high-voltage safety and protection mechanisms. In “Power Electronics Market 2026-2036: Data Centers, Electric Vehicles, and Renewables”, IDTechEx tracks innovations across EV and data center power electronics, drawing deep, cross-application conclusions between EV and data center power electronics.

AI data centers are at a bottleneck and are turning to EV power electronics for inspiration

In November 2022, ChatGPT brought artificial intelligence from science fiction into the real world. Since then, AI has become increasingly mainstream, and leading AI software players, especially across the USA and China, have invested US\$ billions of capital into the development of more complex and powerful AI models. The GPU industry has evolved rapidly to accommodate more advanced training models. Each new generation of server requires more power to be delivered, and incumbent data center rack architecture is simply not able to support the next generation of Nvidia’s “Vera Rubin Ultra” servers.

Rack power levels have exploded from around 20kW pre-ChatGPT to 100kW today, with rack power expected to reach over 1MW by the end of the decade. With incumbent Si power electronics and 480/54V data center architecture, a 1MW rack would require the whole rack space to be dedicated to power conversion and would need over 200kg of copper. Data center efficiency was and still is a key consideration in data center design. Data center power density, the amount of power processed per unit volume, will become an increasingly critical metric for AI data center design.

In “Power Electronics Market 2026-2036: Data Centers, Electric Vehicles, and Renewables”, IDTechEx tracks two key developments in data centers that will facilitate this change: wide bandgap semiconductor adoption, and 800VDC data center architecture. Both of these have been directly enabled by technological innovations in electric vehicles.

800V EV architecture, facilitated by WBGs, enables a step-change in powertrain efficiency

800V EV architecture began to commercialize in 2019 with the Porsche Taycan, and has steadily been adopted across performance and, increasingly, across budget models. The two main drivers for the adoption of an 800V platform relate to increased charging speed and reduced cabling weight and I²R losses, improving overall EV efficiency. The move to 800V has been underpinned by improvements and commercialization of wide bandgap (WBG) semiconductors, SiC in particular.

At 800V, the efficiency improvements of SiC upon incumbent Si IGBT and Si MOSFET technology are significant and outweigh the cost premium of silicon carbide power semiconductors.

With electric vehicles occupying the majority revenue share of the power electronics market, power electronics innovation revolved around EVs and the e-powertrain, for both SiC and GaN technology. This has led to considerable commercialization and cost reduction for both wide bandgap materials, as well as demonstration of their long-term reliability in harsh environments and their material superiority to silicon for high-voltage applications. GaN innovations have been significant due to hype around its potential future adoption into EVs. However, GaN uptake into EVs has been slow, with the material more suited to low-voltage, high-frequency applications.

Wide bandgap and HVDC innovations future-proof data centers for next-generation AI training and inference

Thanks to technological innovations and commercialization across the wide bandgap semiconductor market, largely driven by electric vehicles, the

adoption of wide bandgap technology into data center power electronics has been accelerated. Reference designs for 8kW and 12kW power supply units (PSUs) across leading power electronics players, such as Infineon and Navitas, feature SiC for high-voltage conversion stages and power factor correction (PFC), and GaN for low-voltage conversion stages.


With lower voltage conversion and the possibility for much faster (up to MHz) switching, the material benefits of GaN in data centers are highly desirable; GaN will likely enjoy faster uptake in data centers than in electric vehicles. Switching to WBG semiconductors in data centers leads to efficiency gains, but importantly also leads to significant increases in power density, with smaller SiC and GaN devices, and reduction in the size of passive components, such as capacitors and inductors.

In the same way that WBG semiconductors enabled the 800V shift in EV architecture, they will also facilitate the transition to 800VDC/HVDC AI data centers. With a precedent of 800V architecture, and development of the necessary safety evaluation and processes which can be modified from EV to data center, HVDC data center

architecture development can be accelerated by taking notes from existing EV power architectures. The transition to 800VDC is expected to simplify and future-proof data center architectures. With fewer power conversion stages and lower I²R losses, data center efficiency will increase. Fewer points of failure reduce the chances of downtime. Most importantly, 800VDC data center architecture enables much higher power levels to be delivered to server racks, supporting future GPU generations and AI training models.

The inter-relatedness of power electronics applications complicates the market. Innovations in one application can directly influence large-scale changes in another. In “Power Electronics Market 2026-2036: Data Centers, Electric Vehicles, and Renewables”, cross-industry innovations in materials, manufacture, and device architecture are considered, and clarity is provided on the power electronics industry as a whole.

For more information on this report, including downloadable sample pages, please visit www.IDTechEx.com/PowerElec, or for the full portfolio of related research available from IDTechEx, see www.IDTechEx.com.



➤ AI data center per-rack power requirements are expected to increase over the next five years. To facilitate these changes, key EV power electronics innovations are being brought over the data center industry. Source: IDTechEx

Moore4Power: A new power play for Europe

Infineon is leading a €91 million initiative to advance power electronics, speed development and strengthen Europe's semiconductor resilience.

BY REBECCA POOL, TECHNOLOGY EDITOR

In May, this year, Germany-based semiconductor giant, Infineon, launched a Europe-wide, €91 million project, Moore4Power, to develop next-generation power electronics while strengthening the continent's tech-sovereignty. Over the next three years, 62 partners from 15 nations, including industry leaders, ABB, Airbus, Alstom, Kempower, Signify and Ingeteam, leading universities and research powerhouses, imec, Fraunhofer and CEA, will join forces to advance heterogeneous integration and system-level power electronics.

"We need a new way to manage energy – this is exactly the idea behind Moore4Power," highlights Jochen Koszescha, Senior Director Funding Projects, and Coordination Lead for the Moore4Power project, at Infineon Technologies. "We want to overcome the limits of traditional scaling by bringing together dedicated [power electronics] building blocks to take integration to the next level."

"You know, energy efficiency without power semiconductors just isn't

possible... This is key to the sustainable future of Europe," he adds.

Funding for Moore4Power comes from the participating nations and the Horizon Europe – Chips Joint Undertaking (JU) program that aims to accelerate semiconductor research, pilot lines, and supply-chain resilience across the EU and associated countries, including the UK. The project follows the €73 million, Chips JU-funded project, PowerizeD, which ended last year and was also led by Koszescha with Infineon colleague and Moore4Power Project Manager, Daniela Maier.

PowerizeD had set out to develop digitized, intelligent power electronics for sustainable energy generation, transmission and related applications, including intelligent monitoring and the close integration of gate drivers and control electronic with power devices. As Koszescha points out, project partners spanned the entire development chain from materials research to power modules. Together they developed a highly integrated intelligent power module (IPM) that went beyond conventional designs by integrating the complete gate driver electronics and not just monitoring functions such as temperature sensing. "We also used dual-side reinforced isolation here - this had never been done before and provides big benefits for charging applications in electric vehicles and trucks," he says.

Moore4Power will now take the same Europe-wide organizational



► Launching the Moore4Power project (from left to right): Daniela Maier (Project lead Infineon), Ganesh Chandramouli (Head of Innovation ALSTOM), Jain Chacko (Research Scientist at Fraunhofer ENAS), Joonas Leppanen (Senior Principal Engineer R&D at ABB Finland), Vicky Chatzidogiannaki (Director at INNOVATION DISCO), Jochen Koszescha (Project Coordinator Moore4Power Infineon), Iñigo Polo (General Technology Manager at INGETEAM). [Infineon Technologies]

blueprint. However, this time, the 62 partners - both existing and new - will move beyond module-level design and extend development towards system-level integration, exploring heterogeneous approaches and chiplet-like architectures using hybrid bonding.

“Power electronics are a decisive enabler for energy efficiency and sustainability,” says Koszescha. “With Moore4Power, we are setting the next level of smart integration to achieve significantly higher energy and resource efficiency.”

Heterogeneous focus

For decades, the semiconductor industry has relied on 2D shrinking and monolithic integration to deliver improvements in chip performance, power, area and cost. But as 2D scaling slows, chipmakers worldwide are moving towards heterogeneous design and on-package integration of chiplets and devices to take chip metrics further and accelerate time-to-market. Heterogeneous integration lies at the heart of Moore4Power, with partners getting ready to combine different semiconductor materials, such as silicon, silicon carbide and gallium nitride together with sensing, control and communication function to form tightly integrated systems.

Hybrid bonding is a key technology for advanced logic chips and chiplets, to integrate CPUs and AI processors, and now Infineon and project partners will adapt this manufacturing technology for power electronics. While a logic chiplet, designed to move enormous amounts of data, may contain thousands of fine-pitch copper interconnects, power devices will need far fewer, but much larger and robust copper connections to carry high current in, say, data centre applications.

According to Koszescha, as part of Moore4Power, EV Group, an Austria-based semiconductor equipment manufacturer working on wafer and hybrid bonding, will explore advanced chiplet manufacturing methods, especially hybrid bonding, to heterogeneously integrate silicon and wide bandgap semiconductors within a power module. Power chiplet architectures, designed around module semiconductor building blocks rather than monolithic chips, will enable the development of more scalable, flexible and cost effective systems.

‘World’s largest fab’ for power semiconductors and analog/mixed-signal technologies opens in Dresden

Infineon officially opened its Smart Power Fab in Dresden on July 2, this year, doubling its manufacturing capacities at this site. With an investment of €5 billion from Infineon, the new plant is described as the largest single investment in the firm’s history, directly creating 1,000 jobs.

“Infineon’s new Smart Power Fab sends out a powerful message about Germany and Europe as a center of industry,” said German Federal Chancellor Friedrich Merz at the time. “This investment shows that competitive, cutting-edge semiconductor production in our country is possible.”

“Power semiconductors are a key technology for the energy transition, the future of mobility and our AI infrastructure – meaning they are crucial to make our economy fit for the future,” he added. “The expansion of Infineon’s production capacity in Dresden also strengthens our technological sovereignty and the resilience of important supply chains in Europe.”

According to Infineon, the Smart Power Fab can accelerate production to twice as fast as in the past. The building itself and the ideal machine layout were pre-planned and optimized using a digital twin while system and process clearance is supported by AI algorithms.

The connection to the Infineon plant in Villach as “One Virtual Fab” makes qualification of processes and products significantly faster than in the past. Chips made in Dresden will contribute to the efficient supply of power to AI data centers and are also used in other industrial and automotive applications, such as wind and solar power systems and software-defined vehicles.



► Power Fab in Dresden [Credit: Infineon/Lichtwerke]

As Koszescha says: “We can use silicon for the control electronics and gallium nitride for power switching... Moore4Power is about intelligently combining the strengths of different semiconductor technologies. We’re really looking at smart power.”

Koszescha points to Infineon’s Smart Power Fab in Dresden - also known as Module 4 - that opened in early July (see ‘World’s largest fab...’). As well as securing around €1 billion in funds from the EU chips Act, Infineon has invested more than €5 billion into the facility, described as its largest-ever single investment. Module 4 will combine power semiconductors and analog/ mixed-signal components on 300mm wafers, targeting applications including power supply systems in charging infrastructure, automotive motor control units and data centres.

As Infineon’s Andreas Urschitz, Chief Marketing Officer and Infineon Management Board Member, has said: “With this decisive and courageous move, Infineon is sending a clear message: we are taking risks and investing early to secure scalable, reliable, and resilient supply.”

“We’re taking the ‘fill-the-fab’ approach... at our module four in Dresden,” adds Koszescha. “[Moore4Power] will have the right

technology and devices to take advantage of what we’re doing in Dresden – it’s a real future opportunity.”

Speeding development

As well as integration, Moore4Power is focused on reducing development cycles via AI-assisted models, digital twins and automated workflows. Hardware and software will be designed in parallel to reduce simulation times, with the hope of reducing the time from first fab samples to a validated data-sheet release that engineers can design products around, from several weeks to one week.

“We want to more than halve the time from idea to the real product, which is very, very challenging,” says Koszescha. “We explored this in PowerizeD and now we are going further and using AI and neural networks to simulate power semiconductor performance – we see many possibilities for using AI to speed up product lifecycles, and this is something we’re really working on.”

Heterogeneous integration and faster development cycles aside, Moore4Power is clearly part of a broader European tech sovereignty strategy, in which power electronics are critical. The project is focusing on renewable energy, e-mobility and industrial sectors where power conversion can drive both cost- and energy-savings - and for Koszescha, industry-wide collaboration is

going to be crucial. “It’s not easy having collaboration between more than 60 partners – but we successfully managed this with PowerizeD, and we’re doing it again,” he says.

This time around, Airbus is onboard, which according to Koszescha is a clear indication of Infineon’s central role in Europe’s aerospace-related semiconductor strategy and strengthens European autonomy in aerospace supply chains, including aircraft power distribution. But as Koszescha emphasises, energy, mobility and industry applications, including manufacturing and robotics, all need support from power electronics.

Koszescha also points to the Chips JU €20 million funding opportunity, Call for Reinforcing Europe’s strength in power electronics. This so-called resilience call launched in early July to reinforce Europe’s strength in power electronics, scale power electronics to 300mm device production and adopt agile design.

“Moore4Power is a little bit ahead here as we’re already using some of these ideas,” he says. “Power electronics is key for future sustainability in Europe and we need this funding otherwise we will suffer in the future. All of our partners are ready to support the future resilience of power electronics in Europe.”



➤ With its Smart Power Fab and the Moore4Power project, Infineon is helping to reinforce Europe’s strength in power electronics. [Infineon/mio-motion]

WEBINAR

**YOUR EXPERTISE.
OUR AUDIENCE.
REAL IMPACT.**

A **Power Electronics** webinar helps you share knowledge, generate quality leads and build global brand awareness within the industry.



BRAND AWARENESS

Increase global visibility and credibility



LEAD GENERATION

Generate and nurture qualified leads



THOUGHT LEADERSHIP

Position your experts as industry leaders



OUR 8-WEEK WEBINAR PROMOTION PROGRAMME INCLUDES



Promotion in
Power Electronics
magazine promotion



Targeted exposure
through newsletters to
53,000+ professionals



Promotion in our
active LinkedIn
Group



Featured on the
Power Electronics
website



Moderated by the
editor of *Power Electronics*



LET'S CREATE IMPACT TOGETHER

Our editors work with you to shape the content, engage your audience and deliver measurable results.

Contact: Jackie Cannon

E: jackie.cannon@angelbc.com

T: +44 (0)1923 690 205




**POWER
ELECTRONICS**



Running in the right direction with 48V for AI processor power


As AI processors demand ever-higher power levels, power delivery efficiency is becoming a critical challenge for AI infrastructure. This article explores how 48V Factorized Power Architecture can significantly reduce energy losses, lower operational costs and improve scalability compared with conventional 6V multiphase power delivery approaches in next-generation AI data centres.

BY MAURY WOOD, VP STRATEGIC MARKETING VICOR CORPORATION



➤ Figure 1: Example state-of-the-art AI processor, with CoWoS chiplet package dimensions 55mm by 74mm, or 4,070mm² area

At a certain level of abstraction, modern AI processors simply exchange trillions of bytes of data per second with other AI system processors (typically CPUs), network processors and memory devices, perform quadrillions of parallel vector operations per second, and thermally dissipate two to three thousand watts continuously, around the clock, every day of the year. From that perspective AI processors are the world's most expensive electric heaters. Indeed, a single contemporary AI GPU plus high-bandwidth memory (HBM) accelerator continuously dissipates the same amount of heat as a salon-type hair blow dryer. AI power system engineers are hyper-focused on efficiently supplying the required electrical energy while reducing the thermal dissipation caused by delivering sufficient power to these massive compute engines. The recent shift to an 800V_{DC} based data center rack feed is gaining momentum as the AI infrastructure industry's next move, following the shift away from 12V power delivery networks (PDNs) to more efficient 48V PDNs. However, challenges with voltage regulator current density and current gain at the point of load is causing the industry to head



in the wrong direction. This shift away from efficient 48V PDNs and back to 6V and 2V PDNs requires the management of huge and unnecessary thermal power losses.

A state-of-the-art AI processor (Figure 1) with a thermal design power (TDP) of 3kW, requires 5,000 amps thermal design current (TDC) at 0.6V_{DD}. Note that AI processors can have dozens of supply rails with different voltage and current delivery demands, so the modeling in this article is simplified for clarity. The package dimensions for this chiplet based AI processor example are assumed to be about 55mm by 74mm, or about 4,070mm². This yields a package-level (typically CoWoS or EMIB) power density of about 740mW/mm² and a current density of about 1.23A/mm², supporting the notion of these processors as fierce heaters requiring direct-to-chip liquid cooling to operate continuously and reliably over a period of years.


The proposed industry-standard power delivery architecture: Running in the wrong direction

The legacy standard point-of-load (PoL) AI processor power delivery architecture, whether laterally-placed (LPD) or vertically-placed (VPD), uses interleaved pulse-width modulated (PWM) trans-inductor voltage regulators (TLVRs), typically with dozens of vertically-placed power modules controlled in sequential


“multiphase” outputs. Numerous multiphase VRs are combined to meet the current (both continuous and peak) demands of the various supply rails (multiple logic core rails, HBM memory rails, physical layer SerDes rails, input and output rails and others). In this power architecture, the digital multiphase controller must achieve analog voltage averaging across phases within the processor voltage ripple tolerance, as well as manage accurate output current sharing among the dozens of PoL multiphase power modules.

A key technical characteristic of the multiphase architecture is a powertrain structure that performs voltage transformation followed by voltage regulation. In the context of an AI accelerator card with a 48V input voltage, the first stage of the powertrain, transforms 48V to 12V (4-to-1 fixed-ratio conversion) or more recently and commonly 48V to 6V (8-to-1 fixed-ratio). The second stage, a multiphase TLVR buck regulator, regulates 6V to sub-1V levels, for example 0.6V (10-to-1 buck down) for the processor core rail. The initial transformation stage is commonly referred to as an intermediate bus converter (IBC), and the bus that conducts between stages is commonly referred to as the intermediate bus. Due the very high input current requirements of state-of-the-art AI processors (such as the 5kA TDC at 0.6V used for the analysis in this article), the current across the 6V intermediate bus is 500A. Assuming 200 micro Ohms (μΩ) of PCB power and ground plane lumped resistance across

➤ Figure 2: 6V multiphase powertrain showing transformation then regulation stages and the voltage averaging characteristic



➤ Figure 3: Current industry-standard 6V multiphase power architecture with high intermediate bus power and scalability constraints



► Figure 4: Industry standard multiphase VPD versus Factorized Power Architecture VPD illustrates the 64 times higher intermediate bus thermal dissipation compared to 48V FPA VPD. Both approaches assume 200μΩ bus resistance. Simplified models shown.

this lateral intermediate bus, the I^2R thermal loss is 50W for each AI accelerator card, of which there are dozens in a rack.

The proposed future AI data center architecture is based on 800V power distribution and will rely on separate “sidecar” power racks, creating an intermediate bus, currently proposed to be 6V, to the compute rack processor trays which then feed the multiphase VRs. This power architecture suffers huge intermediate bus PCB thermal losses as the PDN impedance is significantly higher than 200μΩ, and drives the need for silver-plated copper bus bar liquid cooling, and faces critical metal cost and sourcing challenges.

Running in the right direction: 48V Factorized Power Architecture with current multiplication:

The high performance alternative to multiphase TLVR voltage regulation is 48V Factorized Power Architecture (FPA™), using current multiplication at the PoL. This power architecture extends 48V from the AI accelerator input power connector to the AI processor, dramatically reducing thermal losses in AI compute trays and AI factories as compared to the 6V multiphase VR and intermediate bus power architecture.

FPA uses two DC-DC converter stages, but architected as voltage regulation followed by


voltage transformation, the reverse of the standard multiphase VR powertrain architecture. The regulation stage is implemented with a buck-boost or buck DC-DC converter. The transformation stage is implemented with a current multiplier, which divides the regulated 48V factorized bus by a K factor, in this example K = 48, which yields a 1V nominal AI processor supply rail. The inverse of the K factor is the current gain of the voltage transformation stage.

Current multipliers possess unique and innovative high performance characteristics including:

- Ideal DC-DC transformer behavior with no internal energy storage
- High current gain, >40 and high current density >3A/mm² continuous and 6A/mm² peak
- Low output impedance
- High bandwidth to enhance transient response agility
- Best-in-class 8.6A/μs transient response, with 400ns slew rate
- Low EMI with a clean noise spectrum attributable to zero-voltage switching (ZVS) and zero-current switching (ZCS) technology
- Low output voltage ripple
- Module thermal resistance of less than 0.15°C/W
- 1.5mm module height

Running in the right direction, 48V FPA to the PoL dissipates 64 times less heat on the intermediate /

► Figure 5: Intermediate bus thermal power reduction – Vicor 48V FPA versus 12V multiphase, 6V multiphase, and 2V integrated voltage regulation



factorized bus compared to the 6V multiphase VR power architecture, and 576 times less compared to emerging 2V integrated voltage regulator (IVR) solutions.

Using the exemplary state-of-the-art 3kW TDP and 5kA TDC AI processor, assuming a 200 $\mu\Omega$ intermediate / factorized PCB power and ground plane bus resistance, 48V FPA thermally dissipates about 780mW on the factorized bus, while 6V multiphase solutions thermally dissipate about 50W on the intermediate bus.

Aggregate energy and cost savings – 48V to the PoL versus 6V to the PoL at the AI factory level

Summarizing the key point, the 6V multiphase power architecture dissipates 64 times more heat compared to the 48V factorized power architecture. While 800V distribution to AI sidecar power racks is reducing thermal losses, 6V PDNs for multiphase VR power architectures are increasing intermediate bus thermal losses by **64x**. It seems clear the 6V multiphase AI processor power architecture is running in the wrong direction with respect to the pressing industry and broader societal desire to decrease AI factory energy usage.

The bottom line is the bottom line

Frontier AI factories, on the path to 1GW compute capacity, use as many or more than 250,000 GPUs. GPUs are on track to consume 50% of the total utility energy delivered to AI factories of the future. The aggregate 6V multiphase voltage regulator intermediate bus thermal loss for 250,000 GPUs (using the 3kW TDP per GPU model processor) is about 12.5MW (megawatts), while the corresponding aggregate 48V factorized bus thermal loss is about 195kW (kilowatts). This **64x** reduction factor equates to about 12.3MW, or about 108GWh (gigawatt-hours) annually. Again, these processors operate 24 hours a day, 365 days per year. Assuming \$45,000 per 1GWh for wholesale utility-scale electricity, a 109GWh reduction equates to about \$4.9M (million) USD OPEX reduction per year. Over a four-year depreciation period, the aggregate OPEX reduction nears \$20M.

In summary, the \$5M saved annually in a 1GW AI factory electrical OPEX, plus the 55W of reduced heat losses per GPU enabled through the use of 48V to the PoL as compared to the 6V multiphase power architecture, make evident that 48V-based factorized power architecture is running in the right direction, with respect to performance, operational cost, and power savings in the AI infrastructure of the future.



CS INTERNATIONAL CONFERENCE PIC INTERNATIONAL CONFERENCE PE INTERNATIONAL CONFERENCE AP INTERNATIONAL CONFERENCE AIS ANGELTECH INNOVATE SUMMIT

SAVE THE DATE

12-14 APRIL 2027

Sheraton Brussels Airport Hotel

To find out more about our sponsor and speaker opportunities,
contact us today on: +44 (0) 2476 718970

or email:
info@angeltech.net

INDUSTRY SERVICES DIRECTORY

To promote your products and services contact:

Shehzad Munshi T: +44 (0)1923 69015 E: shehzad.munshi@angelbc.com

AIXTRON

Tel: +49 2407 9030 0
Fax: +49 2407 9030 40
Email: info@aixtron.com
www.aixtron.com



Brought to you by:
LUMINUS

Email: apc-e@luminus.com
www.luminus.com/products/APC-SiC-MOSFETs

evatec
process systems

Tel: + 41 81 403 8000
Fax: + 41 81 403 8001
www.evatecnet.com

Greene Tweed

SEALS | CONNECTORS | COMPONENTS

gtweed.com

HITACHI
Inspire the Next

Hitachi Energy

www.hitachienergy.com

WINSOURCE
ELECTRONICS

Tel: + 86-755-83957316
Email: service@win-source.net
www.win-source.net

WÜRTH ELEKTRONIK
MORE THAN YOU EXPECT

Tel: +49 7942 945-0
Email: eiSos@we-online.com
www.we-online.com/en/components/products

Showcase your presence in PEI for 12 months

Promoting your products or your brand through a Corporate Partnership Program in is an economical way to generate interest and drive prospects. Your message will reach over 53,000 professionals worldwide through 3 different mediums of magazine, website and newsletter which creates maximum visibility.

With the largest and most diverse global audience of qualified buyers served by Power Electronics International, secure a corporate partnership and buyers guide entry and receive 12 months of promotion and exposure to the power electronic industry across the magazine portfolio including:

- **Newsletter:** One sponsored newsletter : A sponsor's message (up to 100 words) A 728 x 90 banner displayed at the centre of the alert
- **Magazine:** Listing in 4 issues: A position on the Corporate Partners page for company logo, contact details and web address 300 x 150 pixels.
- **Website:** A button banner 160 x 54 pixels for 12 months plus sponsor logo: 24/7 and 365 days (.jpg or .gif; max. File size: 15k click-thru URL). <https://POWERELECTRONICSMAGAZINE.NET/home>
- All Corporate Partners are included in online buyers guide <https://powerelectronicsmagazine.net/buyers-guide/companies>

WEBINARS

Specialists with 30 year+ pedigree and in-depth knowledge in overlapping sectors

CS COMPOUND SEMICONDUCTOR

SS SILICON SEMICONDUCTOR

PW POWER ELECTRONICS WORLD

PI PHOTONIC INTEGRATED CIRCUITS

SS SENSOR SOLUTIONS

For more information contact:

Jackie Cannon T: 01923 690205 E: jackie@angelwebinar.co.uk W: www.angelwebinar.co.uk
T: +44 (0)2476 718 970 E: info@angelbc.com W: www.angelbc.com

Expertise: Moderators, Markets, 30 Years + Pedigree

Reach: Specialist vertical databases


Branding: Message delivery to high level influencers via various in house established magazines, websites, events and social media

Angel
BUSINESS COMMUNICATIONS



www.EVGroup.com

HIGH-VACUUM BONDING




EVG® ComBond®

Oxide-free, covalent wafer bonding

High-vacuum process environment

Room-temperature process



GET IN TOUCH to discuss your manufacturing needs
www.EVGroup.com