



Electronic Design[®] LIBRARY

Copyright © 2025 by Endeavor Business Media
All rights reserved.

A compendium of articles from *Electronic Design*

In the Age of AI, A NEW PLAYBOOK FOR POWER SUPPLY DESIGN





INSIDE **ELECTRONICS** Electronic Design.

NEW PODCAST FROM *ELECTRONIC DESIGN*: **INSIDE ELECTRONICS**

Subscribe to our new newsletter, *The Weekend Binge*, to get the latest podcast episodes and video content delivered right to your inbox every Saturday.

Join *Electronic Design* and host Alix Paultre, Editor-at-Large, twice a month for the Inside Electronics Podcast. Episodes will feature the latest news, commentary and interviews for the design engineering community. **INSIDE ELECTRONICS** will bring you interesting perspectives and information about the issues, products, technologies and trends impacting the electronics industry and marketplace.

See the latest **Electronic**
podcasts on **Design.**

www.electronicdesign.com/podcasts



In the Age of AI, A NEW PLAYBOOK FOR POWER SUPPLY DESIGN




CONTENTS



02


Introduction



14

CHAPTER 3

Disaggregating Power
in Data Centers



03

CHAPTER 1

Electronic Fuses:
The Future of Power
Protection in Data
Centers?



19

CHAPTER 4


Power Shift: Designing
GaN into Switch-Mode
Power Supplies



09

CHAPTER 2

From 48 V to 800 V:
The Hard Challenges
for High-Voltage DC
Power in Data Centers



26

CHAPTER 5

Scaling AI Data Center
Power Delivery with Si,
SiC, and GaN



In the Age of AI, A New Playbook for Power Supply Design

James Morra, Senior Editor, *Electronic Design*

Can engineers address the AI power crisis from inside the data center? This technical series highlights how power supplies are evolving to handle the current and future demands of AI.

Al is siphoning huge amounts of power from the grid. One way to manage these mounting power demands is to overhaul the power electronics that convert and distribute it around the sprawling data centers driving the AI boom.

For this takeover week, *Electronic Design* consulted with industry insiders and technical experts to fill in the blanks for engineers about how AI's thirst for electricity is transforming server power supplies from the inside out.

At the heart of these data centers are thousands of high-performance AI chips, which are burning through vast amounts of power to train large language models (LLMs) and run other computationally intense workloads. Next-generation GPUs like NVIDIA's Blackwell B100 and B200 consume more than 1,000 W of power each, which is approximately 3X the power budget of a traditional CPU. These new demands are leading to a rapid escalation in data-center power densities, with power-per-rack specifications climbing from 30 to 40 kW to more than 100 kW.

But the processors themselves are not the only culprits behind generative AI's power binge. Compounding the problem are inefficiencies in how power is traditionally converted and distributed inside these colossal data centers.

After electricity enters the rack, it runs through several stages of power electronics before reaching the processor. First, power supply units (PSUs) convert high-voltage AC power into 54 V or 48 V DC before distributing it over busbars to all servers and other hardware in the rack. Next, DC-DC converters inside the server step down the voltage, usually to 12 V, before supplying it over the motherboard to voltage regulators that translate it to the specific voltage used by the SoC. Today, the most advanced chips run on core voltages of approximately 0.8 V. These voltage regula-

tors are the final stage in the power delivery network (PDN), routing current through the copper traces on the PCB and into the pins on the processor's package.

Every one of these conversion steps introduces losses, and the transmission lines between them add even more. These I²R losses—caused by resistance in everything from the busbars and cables in the rack to the copper traces in the server's main circuit board—are increasing as AI chips demand ever-higher currents. As much as 10% to 20% of the power entering a rack is lost as heat before it even reaches the processor. But removing all that heat imposes its own energy cost: approximately 40% of all the electricity used to run a data center is devoted to cooling.

To wring more performance out of every watt, engineers are rethinking power electronics at every stage of the system. However, as rack power demands climb to more than 100 kW, the PSU has become one of the focal points. New power switching technologies such as SiC and GaN bring faster switching, higher efficiency, and better heat dissipation to the table. They're being complemented by more advanced topologies and innovations in everything from gate drivers and digital controllers.

This special coverage dives into the details, highlighting articles from technical experts at Texas Instruments, Infineon, Vicor, Analog Devices, and others. They cover everything from the complexities of high-voltage DC (HVDC) power distribution to the intricacies of power-supply design to advances in circuit protection. Together, they shed more light on what is being done to keep the power flowing to AI data centers.

 [BACK TO TABLE OF CONTENTS](#)




image courtesy of Texas Instruments

CHAPTER 1:

Electronic Fuses: The Future of Power Protection in Data Centers?

JAMES MORRA, Senior Editor, *Electronic Design*

E-fuses integrate power switches, digital controllers, current sensing, and all of the other smarts for hot swapping in a single chip. For TI, they represent the future of circuit protection in AI data centers.

As graphics processing units (GPUs) and other artificial-intelligence (AI) chips push the power limits of data centers, it's becoming vital to prevent the current racing into them from overloading the system and causing costly disruptions.

When a server or other electronic module in the data center fails, it must be “hot swapped” to keep downtime to a minimum. In this case, [hot swapping](#) means removing the faulty hardware and replacing it while the other servers in the rack keep running.

However, that process can introduce a huge amount of inrush current when plugging the server into the rack. The sudden increase in current can stress out the processors, accelerators, and power circuits in the server, potentially leading to component degradation or even failures.

To power up everything safely, a power MOSFET is typically placed in the server with a current sensor and a digital controller. Together, they serve as a circuit breaker for the system, regulating the inrush current when powering up and then keeping track of the supply current during regular operation to prevent short circuits or any other faults. But Robert Taylor, GM of industrial power design at [Texas Instruments](#), said power designers and systems engineers are struggling to scale up these hot-swap solutions within the space constraints of AI servers.

TI and others are rolling out electronic fuses, or eFuses, to safely handle [the huge amounts of power used by AI](#). They integrate the power device, current sensing, and digital control in a single chip that delivers more intelligent [power-path protection](#). Taylor said its latest electronic fuse—the [TPS1685](#)—can detect faults more accurately and intelligently and respond to them faster than other hot-swap solutions. It also includes a black

box for fault logging.

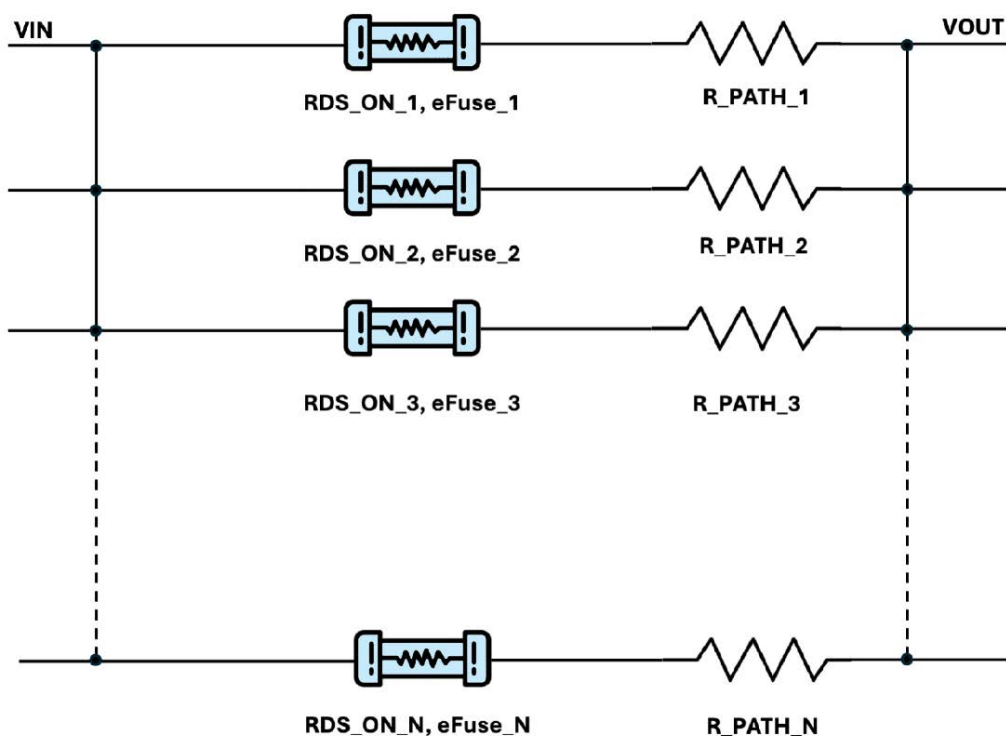
While eFuses are becoming more widely used for hot swapping, TI said the TPS1685 is the first to handle the 48-V DC bus, which is replacing the traditional 12-V architecture as the industry standard for in-the-rack power distribution.

According to Taylor, several of these 20-A eFuses can be stacked together to scale with the rising power demands of AI, which are pushing [power requirements to more than 100 kW per rack](#). “Each one of these chips is able to power about one kilowatt,” he said. “But there is no limit to the number of these that you can put in parallel. The power FETs are integrated, and we control everything between them, so we can make sure that they share current. They are thermally sharing between the different packages as well.”

The MOSFET: The Power Control Switch for Hot Swapping

In a 48-V architecture, the power-supply units (PSUs) in the rack convert the AC used to distribute power around the data center into the 48 V DC used by the servers themselves. The DC power is distributed to the servers through the 48-V backplane that runs down the back of the cabinet. If a server malfunctions, the load is shared by the others to keep everything running long enough for a replacement to be hot swapped into the system.

The capacitors on the circuit board serve to smooth out voltage ripples and remove noise, supplying stable power to the processors, memory, and other building blocks of AI servers. When the server is plugged into the rack, these capacitors require as much current as possible to charge up the processor, accelerator, and other loads, creating a large amount of inrush current for a short time. If the current isn’t limited, it can overload the connectors or the other components in the server or cause sudden fluctuations in voltage



TI said it can safely scale up to higher current loads by stacking the eFuses in parallel and dynamically sharing current between them. (Image credit: TI)

that could reset the servers around it.


The latest high-performance AI chips are burning through more than 1,000 W of power to run AI training and inferencing, which is increasing the amount of current required by AI servers. To power up everything safely during hot swapping, a MOSFET is placed close to the server's power connectors to enable and disable the distribution of power to the processor and other loads. The gate voltage of the power device determines the current passing through to them.

When the power MOSFET is turned off—when the gate voltage (V_{GS}) is under the threshold voltage (V_{TH})—it prevents current flow into the system, blocking inrush current while the server is hot swapped into the system.


But when the power MOSFET is turned on—when the V_{GS} is above the V_{TH} —it enables a constant amount of current to proceed into the system. In this situation, the MOSFET enters the saturation region, where the V_{GS} is the primary factor controlling the drain current. As the V_{GS} increases, the power MOSFET pushes more current into the input capacitor in the system. As the current increases, the

voltage between the drain and source (V_{DS}) of the FET falls, placing the power device in the ohmic region, where the current depends on the resistance ($R_{DS(on)}$) between the drain and source.

These [power FETs](#) require a wide safe operating area (SOA) to prevent the current racing into the system during hot swapping from damaging the power device (or the components around it) or causing it to overheat. To handle larger amounts of current, several of these [power FETs](#) are usually placed in parallel. By spreading out the current over several different power FETs, the heat inside them can also be reduced, said TI. That's also important because high temperatures can increase the overall resistance inside the FET, restricting current.



Bottom view



Top view

Hot-Swap Controller: Protection from the Power MOSFET to the Load

While a shunt or other current sensor checks on the current flowing into the server and the MOSFET connects or disconnects power to the server, the hot-swap controller controls the

eFuses integrate the power switch and all of the smarts for hot swapping in a single device that can be placed above or below the PCB.

(Image credit: TI)

MOSFET and keeps it within the SOA.


The digital controller regulates the gate voltage of the power MOSFET to control the amount of current racing into the system during hot swapping. It's also used to check out the current, voltage, and temperature in the power FET at all other times to prevent [short circuits](#) or other faults. At its heart is a programmable timer, which limits how long the power FET remains in regulation during faults. If the fault condition lasts too long, the power FET shuts down. The power FET must have a large SOA to limit power losses as the timer runs down.

The primary role of the hot-swap controller is to set the system's current limit—it integrates short-circuit and overcurrent protection (OCP) in the event the current runs over the threshold. In many cases, these chips also offer undervoltage lockout (UVLO)—which prevents damage to the power FET that can occur due to fluctuations in the gate voltage—and overvoltage protection (OVP)—which stamps out voltage spikes or problems with the supply voltage. Furthermore, hot-swap controllers often have thermal protection to prevent the FETs from overheating.

Taylor said TI and other semiconductor companies are rolling out hot-swap controllers that integrate gate driving and current-sense amplifiers to more efficiently handle the huge currents used by high-performance AI chips. However, he said they still require separate power FETs to cut power to the load during fault conditions, while [shunt resistors or other](#)

Part doesn't shut off for short load transient (~1ms)

Part shuts off for persistent fault (>2.2ms)



The eFuse is specifically designed to distinguish between the very fast load current steps of AI chips and actual overcurrent faults. (Image credit: TI)

[current sensors](#) are placed with comparators on the same power rail as the FET.

But these solutions are getting more complicated at a time when power boards in data centers are becoming more densely populated. “If you look at the evolution of hot-swap protection, we used to use a lot of separate components—maybe a current-sense amplifier, a comparator, a current sensor, a power FET—and all these different components take up a lot of space on the PCB, and they have made it more challenging for engineers to place all the components” in a way that maximizes safety and minimizes power losses, said Taylor.

He added that given [the growing power demands of AI chips](#), it’s also more challenging to detect fault conditions and shut them down fast before they can damage the processor, accelerator, or components around them.

Inside the First 48-V Electronic Fuse for Data Centers

The future of hot swapping, said Taylor, is the eFuse. By integrating the power switch, digital control, current sensing, [gate driving](#), and all of the other smarts for hot swapping in a single chip, the TPS1685 reduces complexity and saves space in the system.

“We’ve integrated all of these into a single package, so where customers are currently using a hot-swap controller and power FETs today, we can take all of those parts and put them into the same device,” explained Taylor. “We can not only monitor the power FET and make sure it’s always operating in a safe area, but we can also do predictive maintenance on the FET and detect when there are issues with that part of the device.”

At the heart of the new eFuse is a “blanking” timer. It prevents false tripping by enabling the system to distinguish between peak load currents and actual overcurrent faults.

Today, AI server chips have distinct power characteristics, often requiring sudden and large inrushes of current to handle computationally heavy workloads, which can unintentionally cause the circuit breaker to trip. Taylor said the user-programmable timer enables these load transients to travel through the eFuse, preventing them from shutting down the system.

TI noted that the TPS1685 can also be stacked together to handle the rising power demands of AI. Placing discrete power FETs in parallel is challenging due to slight differences in the $R_{DS(on)}$ of the devices and resistances in the PCB traces between them, said Taylor. If any one of the power FETs is forced to handle more current than the others during hot swapping, it can shut everything down even if the system’s total current is under the trip threshold.

To mitigate that, power electronics and systems engineers add margins to make sure that none of the power FETs are taking on more current—and thus, taking on more heat—than the others during hot swapping. The challenge with adding these margins is the complexity of it. “One of the bigger problems with paralleling these devices is having to make sure the mismatches between these different parts are taken into account,” said Taylor.

TI solves the problem by actively sharing the current between the power devices. One of the eFuses is promoted to the role of the primary controller for the system, overseeing the current of the entire system. By measuring the total current instead of monitoring for issues with any one eFuse, the system can assess the situation more accurately during the hot swapping and cut power to the load only when required.

When the controller notices that one of the eFuses is handling significantly more current than the others, it redistributes the current by dynamically increasing its $R_{DS(on)}$. TI

said this dynamic regulation also enables more efficient and robust power distribution by spreading the heat between the devices during peak loads, increasing reliability over the long term.


The dynamic current sharing starts as close as possible to the overcurrent protection threshold of the power FETs to avoid unnecessary power losses at partial loads, according to TI.

The Future of Hot Swapping on High-Voltage DC Power Rails


TI's latest electronic fuse is all about hot swapping on 48-V power architectures. But Taylor said the company is also taking a hard look at high-voltage DC (HVDC) power distribution in data centers.

In response to the rising power demands of AI, Microsoft and other technology giants are trying to relocate the AC-DC power converters in the server rack into a separate disaggregated power rack. The sidecar, as Taylor calls it, would supply power to the server rack at up to 800 V DC rather than AC. Then, a DC-DC converter in the server rack translates 800 V to the 48-V bus that sends power down to the AI processors. There's also the possibility of upgrading the 48-V bus to ± 400 V DC, potentially requiring electronic fuses with high-voltage MOSFETs and robust isolation.

"We have customers that are working up to more than a megawatt of power per rack, so the 400-V architecture is going to be crucial," he said. "We are currently looking at 400-V protection using eFuses as well as different types of architectures for the power-supply units themselves."

to view this article online,  click [here](#)

 **BACK TO TABLE OF CONTENTS**



CHAPTER 2:

dreamstime_Funtap-P

From 48 V to 800 V: The Hard Challenges for High-Voltage DC Power in Data Centers

ROBERT TAYLOR, Sector GM, and SEAN YU, Systems Manager,
Power Design Services and Power Delivery, *Texas Instruments*

Is high-voltage DC (HVDC) the future of power distribution in data centers? If so, power-supply designers will have to work through a lot of challenges to make it a reality.

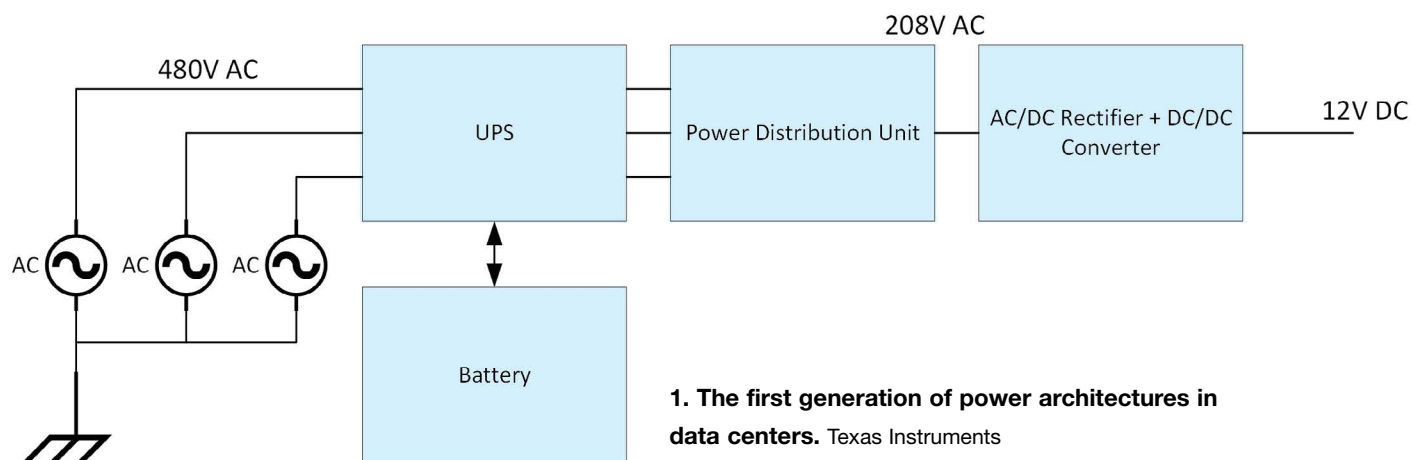
Data centers—specifically those driving the latest innovations in AI—are consuming more and more power. Current research from Goldman Sachs shows that data centers devour 2% of all global electricity today, with expectations for that number to increase to 10% by 2030.

This exposes the shortcoming of traditional power electronics in data centers, from the power-distribution systems supplying AC to columns of servers, to voltage regulators feeding DC power, to [the high-performance AI chips](#) at the heart of it all. The data center's power architecture must undergo major changes to deal with rising power demands of AI, and, in turn, electronic designers will have to solve a number of challenges.

Generational Divide: The Evolution of Power Delivery in Data Centers

To understand where things are headed in terms of power in data centers, it's important to understand the construction of power architectures now. **Figure 1** shows the main building blocks of what you could call a “first-generation” power architecture. In this scenario, three-phase AC (typically a 480-V line-to-line voltage) comes into the data center and feeds an [uninterruptible power supply \(UPS\)](#). The UPS enables battery backup and helps provide a stable AC voltage to the server rack. Inside the server rack, that AC voltage is rectified and stepped down to 12 V using common redundant power supplies in each server blade.

This architecture has been the industry standard for data-center power delivery for decades, and a large portion of systems today still have this configuration. A typical server



rack in this type of architecture will support power ranging from 10 to 15 kW.

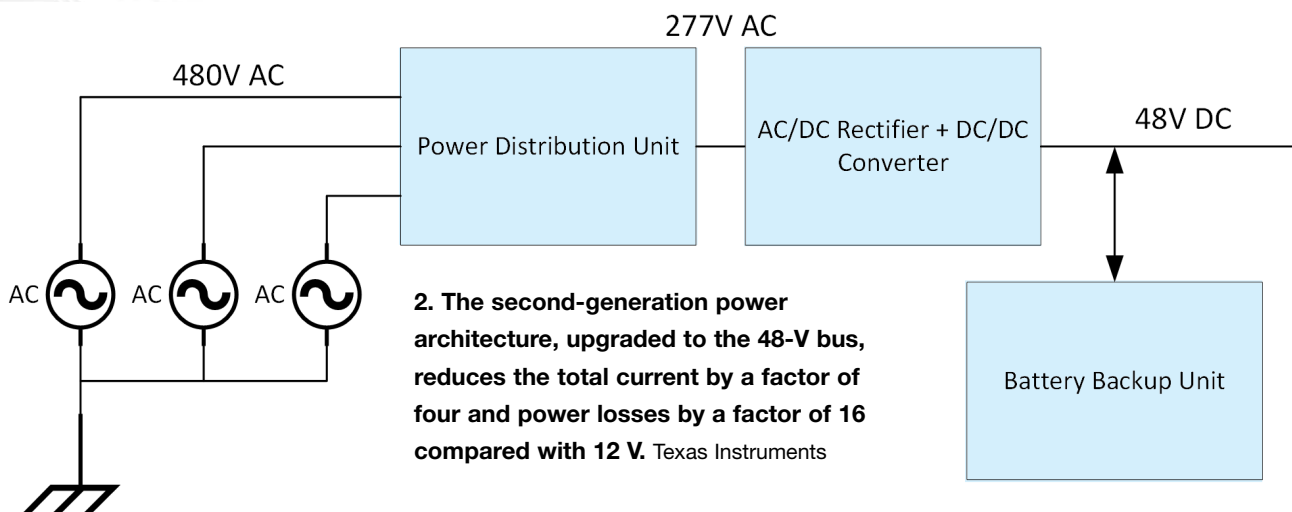
About a decade ago, the technology industry at large started pushing for higher efficiency in data centers. The rise of colossal cloud data centers contributed to increased power levels, which led to the “second-generation” architecture (**Fig. 2**).

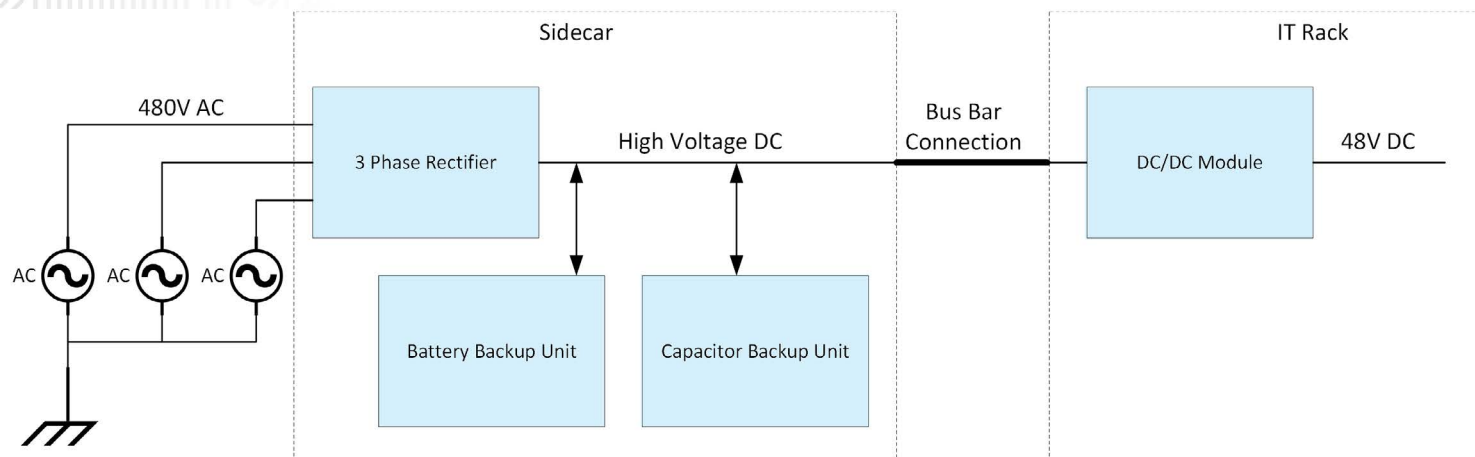
This new system has several differences to the first generation. The output voltage of server power supplies increased to 48 V, along with the consolidation of power supplies into power shelves, also called “open rack” power supplies. A local battery backup unit is incorporated into the rack as well.

All of these improvements lead to a 5% increase in [the efficiency of power-conversion stages](#), while increasing the amount of deliverable power. In a cloud computing data center, the typical rack will have 40 kW to over 100 kW of power.

But as power-hungry AI chips come to dominate data centers, the second-generation architecture is reaching its physical limits. The [power demands of AI data centers](#) are climbing up to [600 kW](#) to 1 MW in a single rack. Today, AI workloads require an immense amount of computing, which necessitates reducing the distance of the physical connections between the graphic processing units (GPUs), the central processing unit (CPU), and networking switches. This configuration means that [big and bulky power supplies](#) need to move out of the IT rack.

That’s why the third-generation architecture introduced the concept of a sidecar, which is basically a separate rack just for power connected to the server rack over [a busbar connection](#) (**Fig. 3**).





3. As AI processors push the power limits of today's data centers, technology companies are relocating AC-DC power-supplies into a separate power rack or sidecar. Texas Instruments

The Design Challenges for High-Voltage DC Power in Data Centers

Now that we've reviewed where the data center of the future is headed, we can discuss the challenges for high-voltage DC power distribution and some potential solutions. Power designers and systems engineers have a long list of questions they need to ask themselves—some of which we may not have even considered yet. But here are several to focus on:

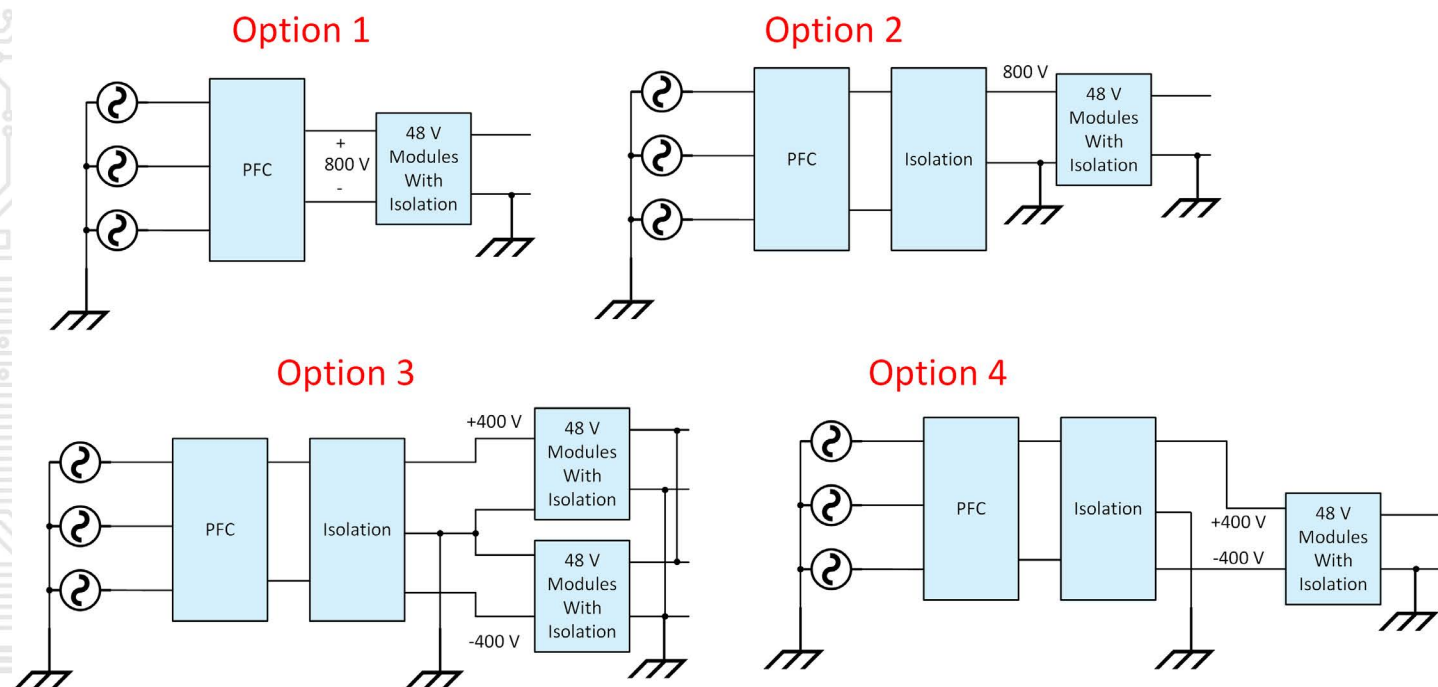
- What is the optimal output voltage level? +400 V, +800 V, ± 400 V?
- What is the role of isolation in the system? Is it strictly necessary on the high-voltage output?
- What are the right power-conversion topologies for the sidecar? What about in the server rack?

If the main motivation for moving the power supplies out of the IT rack is to increase the computing density, why do you have to change the output voltage of the power supplies? The simple answer is related to the busbars that carry the power from the sidecar to the IT rack. If the server rack requires 600 kW of power to run computationally heavy workloads such as [AI training and inferencing](#), it would take 12,500 A of current (not considering any transmission losses) to deliver that power at 48 V.

Because of the current density that's needed, the physical size of that busbar would be very large and weigh close to 200 pounds. These busbars would also require [liquid cooling](#), increasing cost and complexity. Conversely, if you increased the power-distribution voltage to 800 V DC, you would need only 750 A for a 600-kW rack. This current level would allow for air cooling and a weight reduction of 85% per busbar.

The power-distribution voltage must increase—that much is clear. But what is the right voltage level? The +400-, +800-, or ± 400 -V levels are already used [in today's electric vehicles \(EVs\)](#) and the associated charging infrastructure.

The +400-V voltage makes a lot of sense because it's already widely used in today's data centers: [The power-factor-correction \(PFC\) stage](#) in most single-phase AC-DC power supplies outputs +400 V before the LLC stage steps it down to 48 V or 12 V, and electrical components are readily available for use. Engineers also have a solid understanding of 400 V from a safety perspective, as well as regarding creepage and clearance spacing



4. Shown are the different approaches to isolation in a high-voltage DC power architecture. Texas Instruments

in designs. But if the power levels increase more, the busbars could become a problem.

The +800-V level is another viable option for the bus voltage, as it will allow for smaller busbars and higher power-distribution efficiency. This is a relatively new ecosystem for components, though. Engineers will have to work through a range of technical questions regarding safety and spacing. A third option would be to combine the first two options and choose ± 400 V. Its main drawback is that this voltage requires complex [control](#) to ensure balanced loads.

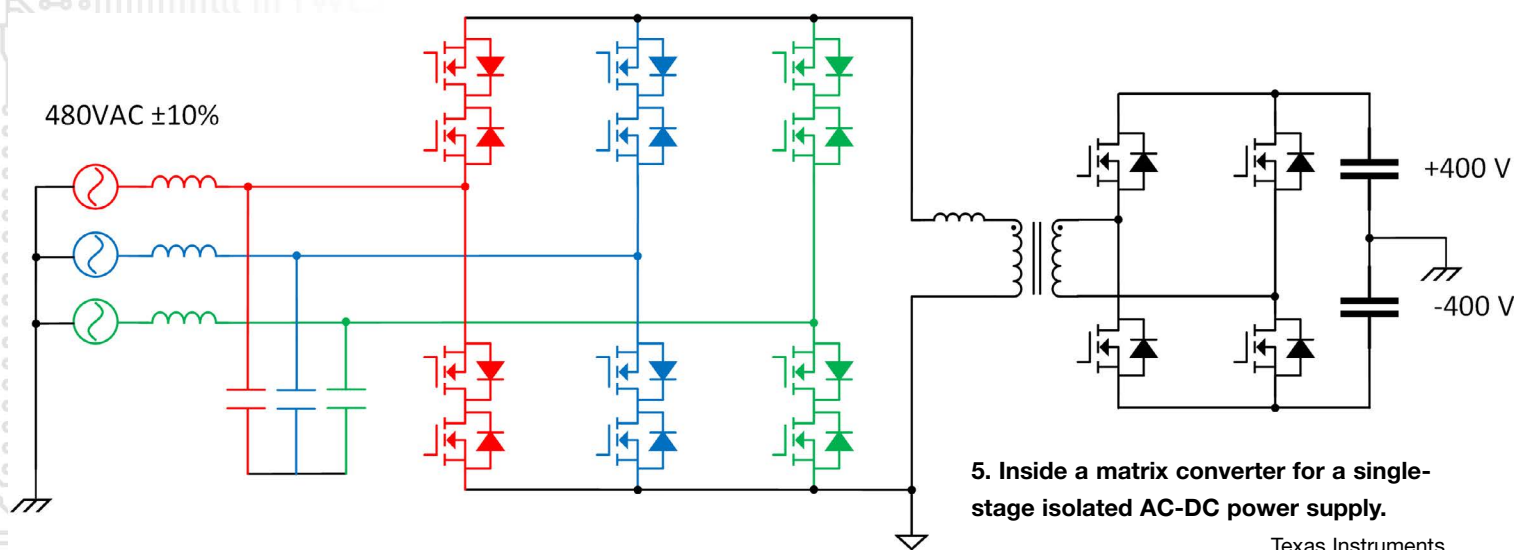
Isolation: How to Handle It in High-Voltage DC Power Systems

It's important to consider all three voltage options, but another issue will influence your choice: isolation.

[Isolation and insulation](#) have two purposes: one is end-user safety, and the other is to keep the ground loops separated. Both are very important for data-center architectures. Because of the power levels needed, it makes sense to have multiple smaller power supplies in parallel, making sure that these power supplies share current.

Figure 4 illustrates several options for the output voltages and isolation schemes. The first option is the most straightforward one, as it involves turning the PFC stage into a separate power-supply unit (PSU). Despite its advantages, there's uncertainty on whether companies will accept it. There are also issues around current sharing and balancing when paralleling multiple non-isolated [AC-DC power supplies](#). Compared to other options with an additional isolation stage, this first option offers the highest efficiency and lowest cost.

The second, third, and fourth options introduce isolation after the AC-DC stage to address current balancing. The third and fourth options use a split rail to generate ± 400 V, with the main difference being the number of busbars required (three versus two). The fourth option needs some additional control to ensure balanced loads on the ± 400 V.



Power Topologies: Balancing Cost, Speed, Efficiency, and Isolation

Another decision involves what topology to use for the AC-DC rectifier. Many factors drive the topology choice, including cost, efficiency, transients on the load, and isolation.

A two-stage approach, where one stage handles the rectifier and a separate stage manages the isolation, is the most traditional and popular way to design a power system. There are many well-known topologies, such as the Vienna rectifier, T-type inverter, or [active neutral point clamp](#), for the rectifier.

Similarly, isolated DC-DC stages, such as three-phase LLC or full-bridge LLC, are available to achieve regulation and isolation. One of the big advantages of this approach is the ability to easily handle energy storage for transients and line-dropout events by adding extra capacitance between the rectifier and isolated DC-DC stage.

Another potential approach is to use a single stage to handle both AC-DC rectification and isolation. This is also known as a matrix converter. **Figure 5** presents a simplified schematic for a single-stage matrix converter.

This type of converter can drive efficiency benefits by reducing the number of switches in the conduction path and by shrinking the overall number of switches and magnetics, thus lowering costs. Some potential drawbacks crop up when it comes to energy storage, though, in addition to concerns about surge voltages.

The matrix converter is also a perfect application for bidirectional switches to help reduce costs and improve efficiency even more. Still, a number of questions and technical details must be figured out to implement this type of design.

The shift to high-voltage DC power distribution is bound to bring about many changes to data-center power supplies. The opportunity to solve complex issues and improve the power supplies is happening now. To meet the power demands of new processors, the data centers of tomorrow will rely on decisions made today.

to view this article online, [click here](#)

[BACK TO TABLE OF CONTENTS](#)



In this Q&A, Vicor contends that ± 400 -V DC power distribution to AI racks in data centers is inevitable.



dreamstime_Studioclever

CHAPTER 3:

Disaggregating Power in Data Centers

JAMES MORRA, Senior Editor, *Electronic Design*,
and MAURY WOOD, VP Strategic Marketing, *Vicor Corporation*

Despite the leaps and bounds in the performance of [the underlying silicon](#), artificial-intelligence (AI) training is pushing the power envelope in data centers. The latest [Stanford AI Index Report](#) shows that the most advanced AI models are becoming bigger, reaching up to 1 trillion parameters and 15 trillion tokens.


As a result, model training is taking more time and resources (up to 100 days and 38 billion petaFLOPS, or PFLOPS), while training costs continue to rise (up to \$192 million). What about the power required to train one of these models? More than 25 million watts.

Tech giants like Amazon, Google, Meta, and Microsoft are turning to nuclear energy to keep up with the colossal amounts of power used to train and run AI. But feeding large amounts of reliable power into their sprawling data centers is only half the battle. The real issue arises inside the server racks themselves, where the power electronics and [the power wires, cables, and busbars](#) are jostling for space with the processing hardware they serve. As power densities rise, managing this internal distribution efficiently is becoming a key issue.

Electronic Design's James Morra discussed the tug-of-war between power and computing with Maury Wood, Vice President of Strategic Marketing at [Vicor](#), who said solving the issue could be as simple as pulling them apart.

How is the underlying architecture of the data center changing to tackle the AI power dilemma?

First and foremost, system designers are making significant efforts to increase compute density, which can be measured in petaFLOPS per liter in EIA-standard 19-in.-wide, or OCP-standard 21-in.-wide, data-center server racks. A single petaFLOP comes out to a



The transition from air-cooled to liquid-cooled rack-scale AI training systems enables a 4X increase in compute density. (Image credit: Vicor)

quadrillion floating-point operations a second.

A related question is, “Why does higher compute density help to reduce the operating expense of training these large AI models?” In short, it’s because inter-processor memory bandwidth and non-optimal latency is a bottleneck on performance. Large model training requires massive amounts of low-latency memory and non-blocking “all-to-all” network fabrics supporting shared access by the dozens of processors within an AI cluster or “superpod.”

Bringing processors, memory, and network switches in closer physical proximity in a rack increases bandwidth and reduces overall inter-processor communications latency, reducing AI model training time. Specifically, the shorter distances defined by a single rack enable the use of passive copper cables instead of active optical transceivers, which are more expensive and power-hungry thanks to the embedded retimers and DSPs.

A typical 800G QSFP-DD and OSFP transceiver consumes about 15 W. Because these supercomputers use tens of thousands of optical transceivers, the power and cost savings of removing all these components are significant—up to 20 kW per rack

What additional steps are being taken to balance compute density with power and cost savings?

The next generation of AI supercomputers has evolved from fan-forced air cooling to liquid cooling. To pose myself another question, “How does this help to increase compute density?” In the previous generation, each eight-processor tray had ten 80-mm fans and a large heatsink that together required eight rack units (RUs) or a compute density of one GPU per rack unit.


The next generation uses direct liquid cooling with low-profile water block cold plates, with two CPUs and four GPUs per one RU tray. This equates to a processor density of four GPUs per rack unit, a 4X increase.

Liquid cooling also eliminates noise and eases the substantial power drawn by the high-RPM 12-V DC fans in these systems. Additionally, by maintaining lower package case and silicon junction temperatures, direct liquid cooling can improve AI processor mean time between failures. It has been reported to be relatively short in air-cooled AI training systems, increasing downtime and operating costs. Higher clock rates are typically also possible in liquid-cooled computer systems compared to air-cooled ones. Both of these outcomes reduce AI model training time and cost.

What else can be done to increase compute density in data centers? What role is power playing?

In previous and contemporary generations of AI server racks, which use three-phase 480-V AC (or sometimes 416-V AC) distribution to the racks, up to 30% of the rack space is consumed with AC-DC rectification, DC-DC conversion to 54 V DC, plus battery backup units (BBUs), capacitor shelves, and/or uninterruptible power supplies (UPS).

To increase compute density and to deal effectively with the prospect of racks that consume up to 140 kW or more, hyperscalers are now advocating an evolution to ± 400 -V DC distribution to next-generation AI supercomputer racks.



Moving the AC-DC rectification and battery-backup (BBU) functions out of the AI training rack contributes to higher compute density. (Image credit: Vicor)

The vision is that the rectification, BBU, and UPS functions are moved out of the 48 RU racks, freeing up space for additional compute and networking trays. This achieves a compute density of 36 CPUs and 72 GPUs, for a total of about 720 petaFLOPS per 48 RUs, assuming rack dimensions of 600 mm width, 1,068 mm depth, and 2,236 mm height. This new system architecture drives up the compute density to about 0.5 petaFLOPS of training performance per liter.

More than anything, the demand for higher AI training performance at lower cost will drive

compute density and subsequently drive the adoption of ± 400 -V DC power distribution.

How does ± 400 -V DC distribution to AI server racks reduce system power and cost?

Existing 480-V AC distribution in data centers typically centralizes the BBU and UPS functions, with large BBU/UPS units supporting multiple AI/ML racks through power distribution units (PDUs).


Because these standalone 2-in-1 units receive AC, they must convert to DC to maintain the battery charge. The BBU/UPS units must also convert the battery output back to AC, and this double conversion process (AC-DC then DC-AC) results in power utilization inefficiency and additional hardware cost. With ± 400 -V DC distribution, no AC-DC rectification function is required at the BBU or UPS.

What are some challenges related to ± 400 -V DC distribution in AI data centers?

The 400-V DC voltage is not safety extra-low-voltage (SELV) level and, thus, presents safety and regulatory issues that must be managed. Also, to preserve the option for 800-V DC powered operation, three conductors (-400 V, GND, $+400$ V) must be run to each rack, which adds cost.

Assuming 140 kW per AI rack, this is 350 A at 400 V DC and 175 A at 800 V DC. Currents as high as 350 A likely require 500 MCM gauge copper cable (380 A ampacity at 75°C), while current of 175 A likely requires 3/0 AWG copper cable (200 A ampacity at 75°C). A 500 MCM gauge copper cable for 400-V DC distribution costs about \$14 per foot, and 3/0 AWG copper cable for 800-V DC distribution costs about \$5 per foot. In large data centers, this almost 3X difference in cable cost is significant.

The cost delta favors 800-V DC distribution, but the 800-V ecosystem is less mature than the 400-V ecosystem, due to use of 400-V DC in EVs. However, automakers are rap-



Vicor's BCM6135 family of power modules supports both 800 V DC and 400 V DC to 54-, 50-, or 48-V DC conversion with high efficiency. (Image credit: Vicor)



idly transitioning to 800-V batteries and DC-DC converters, so the cost issue is dynamic.

One of the biggest challenge areas is handling the high current levels within the rack. Assuming that 400 V DC nominal is converted to 50 V DC nominal using a 1:8 fixed-ratio DC-DC converter, at 140 kW, the conversion yields 2,800 A at 50 V DC. This requires a single silver-plated copper busbar with a cross-sectional area of roughly 1,600 mm² to achieve the required ampacity for air-cooled busbars. A 2.1-meter-long busbar of this cross-sectional area might have a resistance of 5 $\mu\Omega$, and might dissipate up to 45 W at 2,800 A at 20°C, assuming continuous 140-kW rack power.


What are the potential solutions, and how are power electronics playing into the shift?

However, it is possible to liquid-cool the vertical busbar using the existing in-rack liquid-cooling infrastructure and substantially reduce its air-cooled cross-sectional area, up to a factor of 5X (resistance and power dissipation increase with temperature). This represents a significant cost and weight savings.



Liquid cooling of the busbars can also provide tighter control over the maximum voltage drop across the busbar. This reduces the input voltage range of the intermediate bus converters and point-of-load voltage regulation burden on the CPU/GPU accelerator compute modules and network ASIC switch modules. Note that the selection of the 50-V DC connectors also becomes more critical when dealing with thousands of amps of current-carrying capacity to ensure minimal thermal losses.

The OCP Open Rack V3 specification, and the ORv3 HPR (High Power Rack) specification, are industry efforts to address the engineering challenges presented by current and next-generation AI supercomputer power and thermal engineering. Designing next-generation AI supercomputer systems will continue to involve navigating a complex set of engineering and economic tradeoffs.

High-density power modules with low thermal resistance and coplanar surfaces for straightforward mating to liquid-cooling cold plates will play a key role in enabling high-voltage DC distribution to AI supercomputer data-center racks.

to view this article online,  [click here](#)

 **BACK TO TABLE OF CONTENTS**



dreamstime_kyolshin

CHAPTER 4:

Power Shift: Designing GaN into Switch-Mode Power Supplies

BY FREDERIK DOSTAL, *Analog Devices*

Learn more about integrating GaN power FETs into switch-mode power supplies and what you must change when upgrading from silicon MOSFETs.

Silicon is the most amazing material for electronics. The ability to grow pure bulk silicon and dope for both p- and n-type properties has been a huge boost for the power electronics industry, leading to low-cost, high-performance switches that permeate virtually everything with a battery or a plug.

As a result, designers' experience with building circuits out of silicon is enormous. Their familiarity with it has enabled the industry to push the boundaries of silicon further and further over time. However, while silicon performs well in various applications, certain material traits limit advances in speed, power density, and temperature range—all of which matter to the latest [switch-mode power supplies](#) (SMPS). The pressure to seriously consider an alternative technology is increasing the focus on silicon carbide (SiC) and gallium nitride (GaN).

[SiC](#) has been used in electronics for a very long time. Early applications involved light-emitting diodes (LEDs). Lately, it's being used as a power-stage component in power supplies due to its high-temperature and high-voltage capability. Switches and diodes with voltage ranges way above 1,000 V are available.

One other technology that can replace or enhance silicon circuits in power applications is GaN. Widely used for consumer [fast-chargers and power adapters](#), it's also now gaining relevance in areas such as electric-vehicle (EV) onboard chargers (OBCs) and [DC-DC converters](#), where it's necessary to handle thousands of watts at a time safely and robustly. GaN also has the potential to alleviate some of the power challenges posed by AI in data centers, weaving its way into [the power-supply units \(PSUs\) in server racks](#).

GaN is becoming more and more popular for these switch-mode power supplies. For circuit designers interested in using this relatively new technology, it's necessary to not only understand the benefits and the challenges, but also gain experience with it.

The Power-Handling Properties of GaN

GaN was used for the first time in power switches as a replacement for silicon field-effect transistors (FETs) in SMPS in 2012. These prototype pGaN HEMTs showed higher power-conversion efficiency compared to a standard silicon FET device.

The main difficulty with GaN power technology was—and still is—driving down its cost. The ability to grow large single crystals for manufacturing large, high-quality wafers full of GaN power devices remains a challenge.

Working through these challenges will be worth it, though. GaN comes with a wide range of advantages over silicon when it comes to converting power. The main benefits of GaN power devices are the lower drain and gate capacitance for a given current and voltage rating. In addition, GaN switches are physically smaller than silicon, resulting in a more compact solution.

The material properties of GaN mean that it also has a high breakdown voltage, which is useful in applications running at voltages of 100 V and above. However, below 100 V, the power density and the capability of fast switching can give this technology advantages, such as higher power-conversion efficiency, when designing different power supplies.

GaN is a [wide-bandgap \(WBG\) semiconductor](#), which means that the bandgap voltage is 3.4 eV versus the 1.1 eV of silicon. However, figures of merit matter differently in power-supply design. A valuable use case would be in 400-V intermediate bus applications, such as in 240-V AC power converters where we use 650-V breakdown voltage FETs with a drain source current of roughly 30 A. This system requires a gate charge of 93 nC when using a silicon FET, versus only 9 nC using a GaN FET. An application utilizing such switches would run in power levels roughly between 1 and 8 kW.

The benefit of using the GaN device with the small gate capacitance results in much faster switch transition times and reduced switching losses. Ultimately, it would lead to higher power-conversion efficiency, especially at higher switching frequencies with smaller magnetics.

The Unique Challenges of Using GaN in Switch-Mode Power Supplies

A number of challenges emerge when replacing [silicon MOSFETs](#) with GaN power devices. These challenges relate to the requirements for gate drive, fast-changing voltages during switching, and high conduction loss during dead times.

First, GaN switches typically have lower gate voltage ratings than silicon FETs. Most manufacturers of GaN devices recommend a typical gate-drive voltage of 5 V. At the same time, it's not unusual to have GaN devices with an absolute maximum rating of 6 V, which doesn't leave much headroom between the recommended gate-drive voltage and the critical threshold, above which the voltage can damage the device. This limitation, along with the fact that the gate charge in GaN devices is so small, means that driver stages must strictly limit the maximum gate-drive voltage to avoid damage to the GaN device.

Second, one must deal with the fast-changing voltages—also called dv/dt —of the power supply's switch node. These transient voltages may cause false turn-on of the bottom switch. The gate of a GaN device is relatively small. As a result, any fast voltage changes in the vicinity, such as the switch node, may capacitively couple onto the small gate of the GaN switch and turn it on. Gaining more control of the turn-on and turn-off profiles requires a separate pull-up and pull-down pin and a carefully designed [PCB layout](#).

Lastly, GaN FETs have a higher conduction loss during dead times. These are the times when both the high-side and low-side switch of a bridge configuration are turned off. Dead

times are necessary to prevent a short circuit from the high-side voltage rail to ground. During the dead time, the low-side switch typically develops current flow through a body diode of the low-side switch, which is what causes the conduction losses.

One way to solve this problem is to strictly minimize the length of these dead times. This needs to be done without generating overlapping times of the high-side and low-side switches, which can cause a short circuit to ground.

One other point to mention is the fact that GaN offers a wider conversion range. The fast rise and fall times provide a smaller duty cycle than with silicon MOSFETs.

Switching It Up: Upgrading From Silicon MOSFETs to GaN Power Devices

For many years, silicon has formed the backbone of the power-conversion industry. Now that GaN switches are available for power-supply designers, the question is, “Are they simply a drop-in replacement or do you have to redesign the power stage around them?”

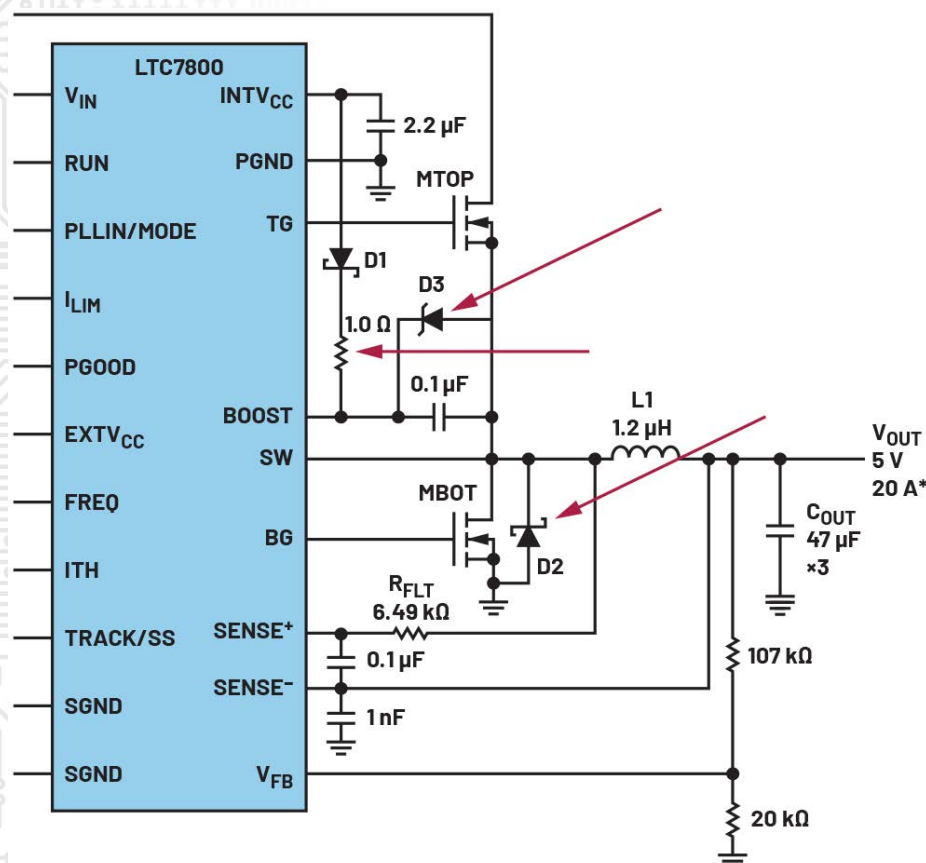
Figure 1 shows a power stage in a typical buck-regulator SMPS. The red arrows indicate additional components that may be necessary when using GaN switches in a SMPS.

Unlike silicon MOSFETs, GaN switches don’t have a body diode—they have a different mechanism to yield similar results. Only majority carriers are involved in GaN device conduction, so there’s zero reverse-recovery charge (Q_{rr}). However, the GaN FET doesn’t

have the forward voltage of the body diode, as with silicon MOSFETs, so the voltage across the GaN FET may get quite large. Thus, the power losses during the dead time are quite high. This is why it’s important to reduce the dead time when using GaN switches compared to silicon switches.

Power designs use a silicon MOSFET’s body diode excessively during the dead time in a SMPS. In a buck regulator’s low-side switch, the current flow through its body diode provides the continuous current flow demanded by the inductor. Without a body diode in the low-side switch, every bit of dead time would cause the switch node in a buck regulator to go to minus infinity voltage. Most certainly the circuit would lose energy and eventually blow up due to voltages outside of the rated voltage of the switch, before reaching minus infinity.

If the source and gate are at the same potential when using a GaN switch, but with a continuous current source such as the inductor, the GaN FET will turn on in reverse.



1. Necessary components to consider when using GaN technology as power switches in a power stage of the [LTC7800](#) buck converter. Analog Devices

Since GaN switches don't include a p-n junction body diode, the low-side switch needs to be constructed with an alternate current path around the low-side switch, allowing for current flow during the dead time. *Figure 1* shows a simple Schottky diode (D2) placed between drain and source of the low-side GaN switch. This diode will quickly take over the inductor current flow during the dead time of the circuit.

During the reverse conduction, the drain and the source get flipped due to the symmetry of the GaN power FETs. The gate remains at ground potential, but the switch node is self-biased to be the minimum turn-on threshold of the GaN FET. This low voltage is the minimum threshold needed to turn the GaN FET on (typically GND-2V to GND-3V). Since the V_{GS} isn't optimized, the R_{ON} suffers during reverse conduction. The external Schottky is an alternative path without turning the GaN FET on in the reverse conduction.

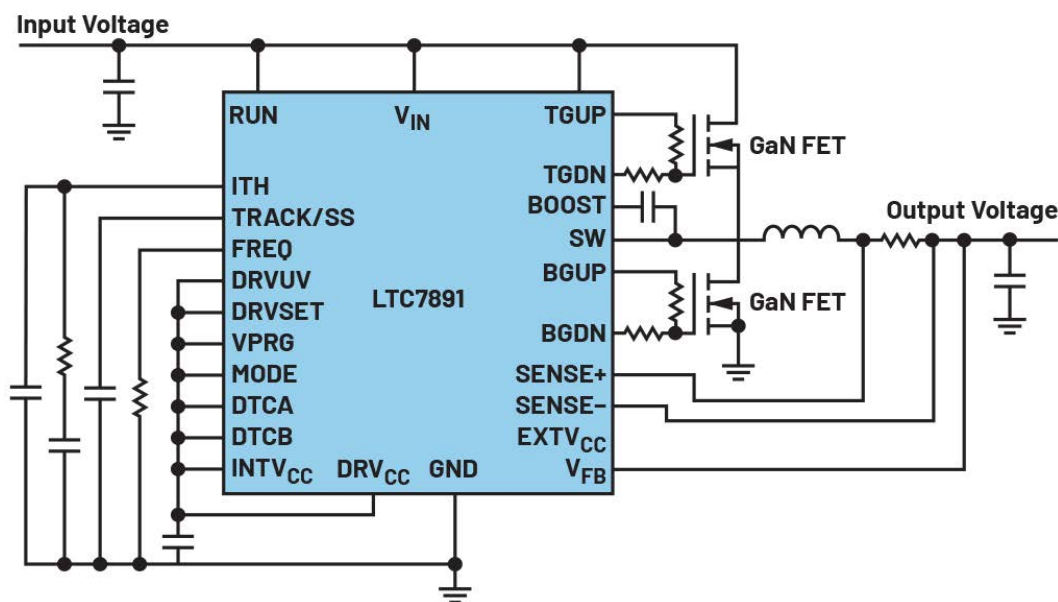
The second modification to the circuit when using GaN switches is the resistor in series with diode D1 (**Fig. 2**). The resistor is used to supply the voltage to the high-side driver of the circuit coming from the INTVCC supply voltage. The resistor may also be needed to limit peak currents for the high-side driver.

Lastly, the Zener diode (D3) may be needed to prevent voltage spikes from becoming excessive on the high-side driver voltage supply.

While the additional components in *Figure 1* look relatively simple and straightforward, ensuring that such a circuit will run reliably in all operating conditions, they require fine-tuning and thorough evaluations on the test bench. Also, variations of component values over production and over aging will need to be considered. The ultimate risk is permanent damage to the GaN switches.

How to Choose Switching Controllers and Gate-Driver ICs for GaN

One way to avoid the critical evaluation process of protection functions in a GaN-based power stage of the SMPS is to use a power-supply controller IC specifically designed for GaN. Selecting a dedicated controller makes a GaN power-supply design simple and robust. All the challenges mentioned earlier are addressed and solved with such controllers. *Figure 1* shows the simplicity of a step-down power design using GaN FETs con-



2. A dedicated GaN controller yields a robust and dense power-supply circuit.

Analog Devices

trolled by the [LTC7891](#), a single-phase buck controller designed for GaN FETs.

These switching controllers also offer the flexibility needed to work with the different GaN switches on the market today. GaN power technology is far from having finished its path of development and innovation. Though future GaN switches will bring better performance to the table, they may require slightly different handling compared to the switches that are readily available today.

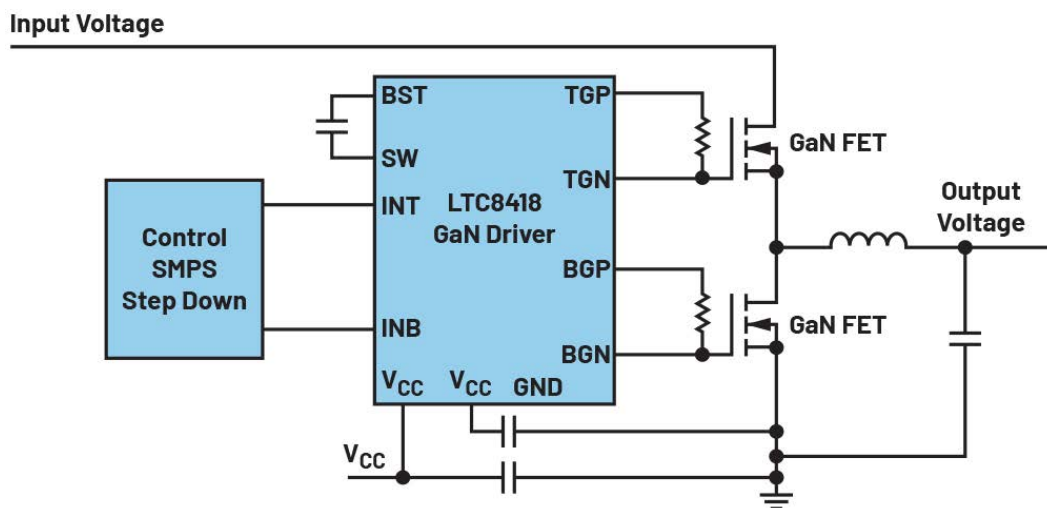
Devices such as the LTC7891 in *Figure 2* offer a dedicated up and down gate-drive pin for both switches. With this, the rising and falling slope of the gate voltage of the GaN switches can be controlled separately. It allows for driving the power stage with GaN switches that have minimal ringing and overshoots.

Compared to a standard silicon MOSFET buck controller, the LTC7891 also features separate gate-drive pins for rising and falling edges (*Fig. 2, again*). However, many other differences separate the LTC7891 from existing controllers designed for silicon switches. There's an internal bootstrap switch to prevent the overcharging of the high-side driver during dead times. It's implemented reliably without the need for external components.

Another important feature is very fast dead-time control. It leads to reliable operation and yields significant power-conversion efficiency improvements, while also allowing for high switching frequencies. The LTC7891 is rated for switching frequencies of up to 3 MHz. One other unique feature is the possibility of accurately adjusting the gate-drive voltage from 4 to 5.5 V. This optimizes the V_{GS} that's needed for various GaN FETs available in the market.

Using passive components with a standard silicon controller, or implementing a dedicated GaN controller, are both viable options. But engineers can also consider using a standard controller IC and adopting a driver stage optimized for GaN. This can also take care of the challenges with GaN and allows for a simple and robust design.

Figure 3 shows the power stage of a buck regulator implemented with the [LT8418](#) driver IC. The driver comes in a very small wafer-level chip-scale package (WLCS) that enables very low parasitic resistances and inductances for low-voltage offsets resulting from fast current changes.



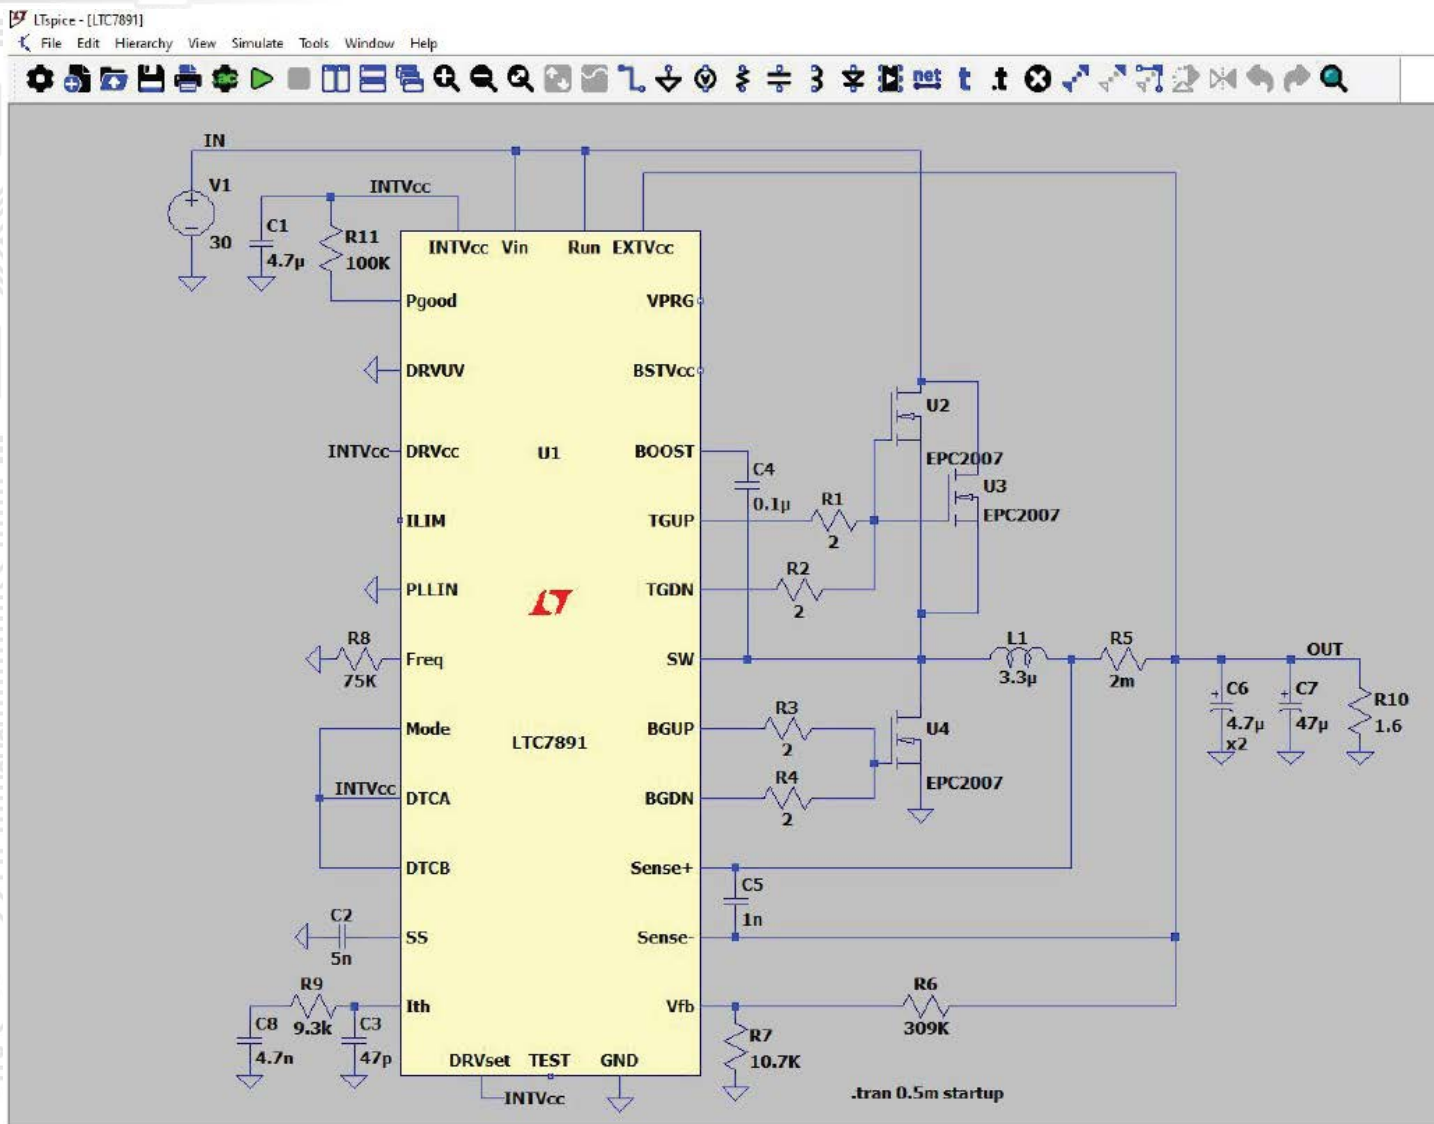
3. A dedicated GaN driver controls a power stage based on logic PWM signals from a heritage silicon MOSFET controller. Analog Devices

Once suitable hardware, controller ICs, and GaN switches have been selected, a great way to get first evaluation results is to use a detailed circuit simulation. LTspice from Analog Devices offers complete circuit models that may be used for simulation. It offers a convenient way to learn about using GaN switches. *Figure 4* shows a simulation schematic of the LTC7891. A dual-channel version, the [LTC7890](#), is also available.

While GaN technology is brilliant for building FETs and using them in an advanced power stage, GaN isn't necessarily able, nor cost-effective enough, to be used in the control circuitry for a switch-mode power supply. Therefore, we will see a hybrid approach for the foreseeable future.

The controller will be silicon-based with highly optimized control and drive circuitry to drive a high-power GaN switch. This approach is technologically available today and it's cost competitive. However, it will require the use of multiple die in one circuit. This entails either having the GaN switches separate, as shown in this article, or by incorporating multiple die in a fully integrated power-converter IC or a μ Module power-supply solution that integrates many of the passive components, including the inductor.

4. LTspice can be a useful simulation tool for GaN power supplies. Analog Devices



As noted, growing large, high-quality GaN wafers remains a challenge. Since around 2010, the mainstream choice in GaN manufacturing is high electron mobility transistors (HEMTs) on silicon due to the larger possible wafer diameters and lower cost associated with existing silicon processing infrastructure.

Early technical challenges with this approach have been resolved. Still, the technology requires years of further development. Instead of being grown as bulk crystals like silicon and SiC, GaN devices are built using GaN epitaxy on silicon wafers. GaN-on-diamond is one potential way of processing GaN switches in the future.

What's the Future for GaN Power Technology?

GaN has reached a solid point in its development where many SMPS can be designed with it. However, there will be further development with each new generation of GaN switches. The existing SMPS controllers and drivers for GaN from ADI are a flexible and dependable way to work with GaN FETs from different vendors now and in the future.


We're actively progressing toward broader adoption of GaN technology. While GaN switches by themselves are quite robust today, it's evident that more time and development will be necessary for engineers to accept the reliability of the switches. The manufacturing processes of GaN switches will also improve further, increasing yield and reducing defect density, thus reducing the cost and enhancing the reliability of GaN switches.

Expect more and more specific GaN drivers, such as the LT8418, and switching controllers like the LTC7890 and LTC7891 buck controllers, to arrive to simplify the implementation of a GaN-based SMPS.



Currently, the most common GaN device voltage ratings are 100 V and 650 V, which is why the first power supplies using GaN operate within these ranges. But the unique characteristics of GaN—specifically, its small gate charge—can also scale down to lower voltages.

In the future, we will also see GaN used in power supplies with maximum voltages as low as 40 V. At the other end of the spectrum, GaN switches may also scale up to 1,000 V. At such high voltages, the fast-switching speeds of GaN can be a difference maker.

Semiconductor materials that expand the operating range and power density of power supplies will continue to be developed. Silicon was the start of it all, and it's still the center of the power electronics industry. But GaN is going to play a bigger role in the next 10 to 15 years, too. Everything from EVs to AI data centers requires more power, higher density, greater robustness, and higher efficiency. GaN provides the opportunity to keep pace with these innovations.

to view this article online,  [click here](#)

 **BACK TO TABLE OF CONTENTS**



dreamstime_kyolshin

CHAPTER 5:

Scaling AI Data Center Power Delivery with Si, SiC, and GaN

BY JUSTIN CHOU, Application Marketing Director, *Infineon Technologies*

Learn more about
how SiC and GaN are
redefining power-supply
design to meet the
growing demands of AI
SoCs.

Large language models (LLMs) and other neural networks draw substantial power when processing complex artificial-intelligence (AI) and machine-learning (ML) workloads. Designed for traditional server configurations, conventional power-supply units (PSUs) can't efficiently keep pace with the demands of GPU-based AI accelerators. To meet evolving workload requirements, data center operators need power-delivery systems that scale within existing thermal and physical limits—without increasing cost, power use, or cooling overhead.


The PSU is the power backbone of the data center, converting high-voltage AC power from the grid into the lower-voltage DC used by all of the components within the server rack. To deliver higher output and improved efficiency while minimizing losses and managing heat, many engineers are combining silicon with wide-bandgap (WBG) semiconductor materials such as silicon carbide (SiC) and gallium nitride (GaN), which improve switching performance and thermal stability under heavy AI workloads.

This article outlines the complementary roles of all three power semiconductors, focusing on how these technologies have been integrated over several generations of Infineon's PSU reference designs ranging from 3 to 12 kW.

Silicon: Can It Meet the Power Demands of AI?

While silicon has long been the foundation of power electronics, its physical limitations are increasingly apparent in high-performance, high-density applications such as AI server racks and other data center infrastructure.

[High-performance GPUs](#) now operate at up to 1,000 W each, with projections reaching 2,000 W and beyond by the end of the decade—comparable to the total draw of a legacy server just a few years ago. As AI hardware and applications continue to scale, PSU power ratings must increase from typical levels of around 800 W to 5.5 kW and higher. Infineon



1. The current roadmap for data-center power supplies ranges from 3 kW up to 12 kW.

Infineon Technologies

estimates that data centers could consume up to 7% of global electricity by 2030—roughly equivalent to India's current national energy use.

This surge in power demand—combined with grid constraints such as limited substation capacity, transmission bottlenecks, and variable renewable generation—highlights the urgency of deploying more efficient power systems into data centers. Rather than absorbing the steep cost of new construction or [shifting AC-DC power supplies into separate racks](#), many operators are opting to increase power density while staying with existing rack footprints.

As demand grows for higher voltages, faster switching frequencies, and greater power density, [conventional silicon-based devices](#)—such as metal-oxide-semiconductor field-effect transistors (MOSFETs)—are reaching their performance and efficiency ceiling.

To support the requirements of compute-intensive AI and ML workloads, the latest PSUs leverage [WBG semiconductors](#) such as SiC and GaN with lower on-resistance ($R_{DS(on)}$), higher switching speeds, and reduced losses. While these materials have superior power-handling properties, they're also edging closer to silicon in cost. As mass production scales up and design processes mature, prices are set to drop substantially. These trends position SiC and GaN power devices as practical complements to silicon in the design of dense, efficient, and thermally optimized PSUs for the AI era.

But there are tradeoffs with all three of these technologies when it comes to efficiency, power density, and cost. To understand where each technology fits into the present and future landscape of AI power delivery, we'll take a closer look at how Infineon's power-supply designs have evolved to address the ever-increasing power demands of AI data centers. Solutions scale from 3 to 12 kW and output voltages reach up to 50 V DC to support [busbar distribution](#) in AI server racks (**Fig. 1**).

These designs deliver efficiencies of up to 98%, meet strict hold-up time requirements, and achieve power densities as high as 100 W/in³. Several of them align with Open Compute Project (OCP) ORv3 standards and support rack-level loads of up to 250 kW and beyond.



The 3-kW PSU: Benefits of SiC-Based Power Switching

The [first-generation AI PSU](#) features a front-end AC-DC bridgeless totem-pole power-factor-correction (PFC) converter followed by a back-end, isolated, half-bridge, series-parallel resonant LLC converter. The 3-kW PSU leverages [CoolSiC MOSFETs](#) and CoolMOS switches to achieve a peak efficiency of 97.5% with an internal DC fan and 97.4% without one. Its compact form factor— $73.5 \times 520 \times 40$ mm—yields a power density of approximately 32 W/in³.

In the OCP ORv3 architecture, each power “shelf” in the server rack houses several PSUs. While the input for the power shelf is three-phase AC ranging from 400 to 480 V, the input for each PSU is single-phase AC from 230 to 277 V. The PSU outputs a tightly regulated DC voltage such as 48 or 50 V to the busbar that runs down the server rack, feeding power into the AI servers and the battery backup units (BBUs) located under the rack.

At 50% load, the main losses in the totem-pole PFC stage originate from the half-bridge boosting switches, which operate predominantly in hard commutation. To address this, the design uses 650-V CoolSiC MOSFETs in 4-pin packages, which feature low parasitic capacitances to support high switching frequencies. In the slow leg of the totem pole, a silicon CoolMOS replaces the diode rectifier to enable [synchronous rectification \(SR\)](#), which reduces conduction losses.

[SiC MOSFETs](#) excel in high-current, high-voltage applications. They have minimal reverse-recovery charge (Q_{rr}) and a more temperature-stable $R_{DS(on)}$ than silicon MOSFETs due to the higher thermal conductivity of SiC. They optimize performance in both soft- and hard-switching topologies. In the totem-pole PFC stage, which operates in a [hard-switched, continuous-conduction mode \(CCM\)](#) when managing kilowatts of power, SiC supports higher switching frequencies, reduced commutation losses, and improved overall efficiency.

In the LLC converter stage, transformer losses are a primary contributor, along with significant switching and driving losses. To optimize performance and power density, the converter operates at a relatively low resonant frequency of 93 kHz. Silicon CoolMOS devices offer the best cost-performance ratio for this frequency range and board size.

Complementary Properties of Silicon, SiC, and GaN

The second generation is a [high-frequency, high-density 3.3-kW PSU](#) that leverages all three power switching technologies to maximize the efficiency, power density, and overall cost of the system. The front-end AC-DC converter features a two-phase, interleaved, bridgeless totem-pole PFC with CoolSiC, while the back-end isolated DC-DC stage implements a 500-kHz, GaN-based, half-bridge LLC converter with full-bridge rectification.

The PFC stage comprises a pair of interleaved totem poles featuring a total of four 650-V CoolSiC MOSFETs for the fast legs. [EMI performance](#) is also improved through careful circuit refinement. Notably, substantial reductions in both line and neutral positive-peak EMI are achieved by optimizing the zero-crossing sequence in the PFC and using a total of four 600-V CoolMOS SJ MOSFETs in the slow leg of the bridgeless totem-pole PFC.

Compared to the first-generation 3-kW PSU, the 3.3-kW design operates at a significantly higher resonant switching frequency of 500 kHz. To support this, the LLC stage uses a total of four CoolGaN devices in place of CoolMOS, optimizing efficiency and power density at high switching frequency.

GaN power devices thrive in situations where high-frequency switching is important.

Their lack of a body diode results in zero Q_{rr} , which minimizes switching losses. This enables shorter dead times, thus reducing unnecessary conduction losses. GaN also features very low output capacitor charge (Q_{oss}), which helps facilitate [zero voltage switching \(ZVS\)](#), especially in soft-switching LLC converters. It can also enable more efficient operation in hard-switched, half-bridge topologies such as totem-pole PFC.

The design also features Kelvin source connections. The 80-V OptiMOS switches used for synchronous rectification in the LLC converter have dedicated gate drivers, further minimizing losses and enabling accurate, high switching frequencies.

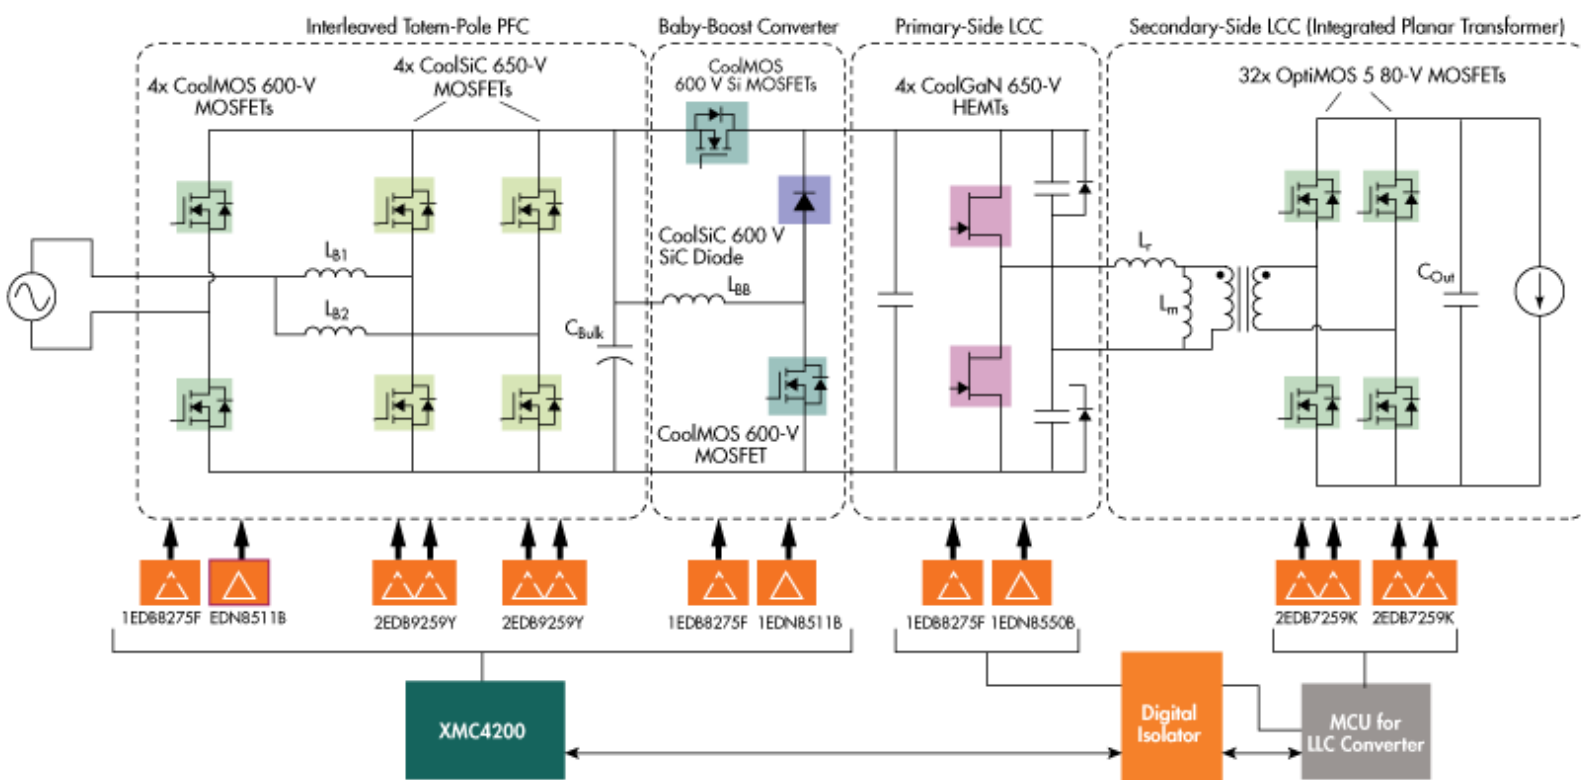
The design incorporates a “baby-boost” converter in between the AC-DC rectifier and the DC-DC converter stages to meet the 10 ms of hold-up time required in server applications—without significantly increasing bulk capacitance. This addition reduces the capacitor volume, contributing to higher overall power density (**Fig. 2**).

The PSU delivers a peak efficiency of 97.4% with any internal cooling fans and a power density of 98 W/in³, based on compact outer dimensions of 72 × 192 × 40 mm.

The 8-kW PSU: More SiC, More GaN, and More Silicon Power Switching

As high-performance AI chips drive up power-per-rack specifications in data centers to more than 100 kW, supplying more power in a smaller space is becoming critical. Infineon introduced a single-phase [8-kW PSU reference board](#) to manage these mounting power demands. This design builds on the architecture of the 3.3-kW design and scales it up to higher power levels, leveraging all three semiconductor technologies to an even greater degree to optimize performance, size, and cost.

The front end comprises a bridgeless, interleaved, totem-pole PFC converter that regu-



2. Shown is a simplified schematic of the 3.3-kW reference design, the REF_3K3W_HFHD_PSU. Infineon Technologies

lates the intermediate high-voltage bulk. Due to elevated current stresses, all of the power switches in the PFC stage—in both the fast and slow legs—are based on SiC. Each of the interleaved totem poles have a pair of SiC MOSFETs placed in parallel to achieve a lower equivalent $R_{DS(on)}$ and reduce conduction losses. The fast legs feature a total of eight 650-V, 40-m Ω SiC MOSFETs, while the slow legs leverage a total of four 650-V, 10-m Ω SiC MOSFETs.

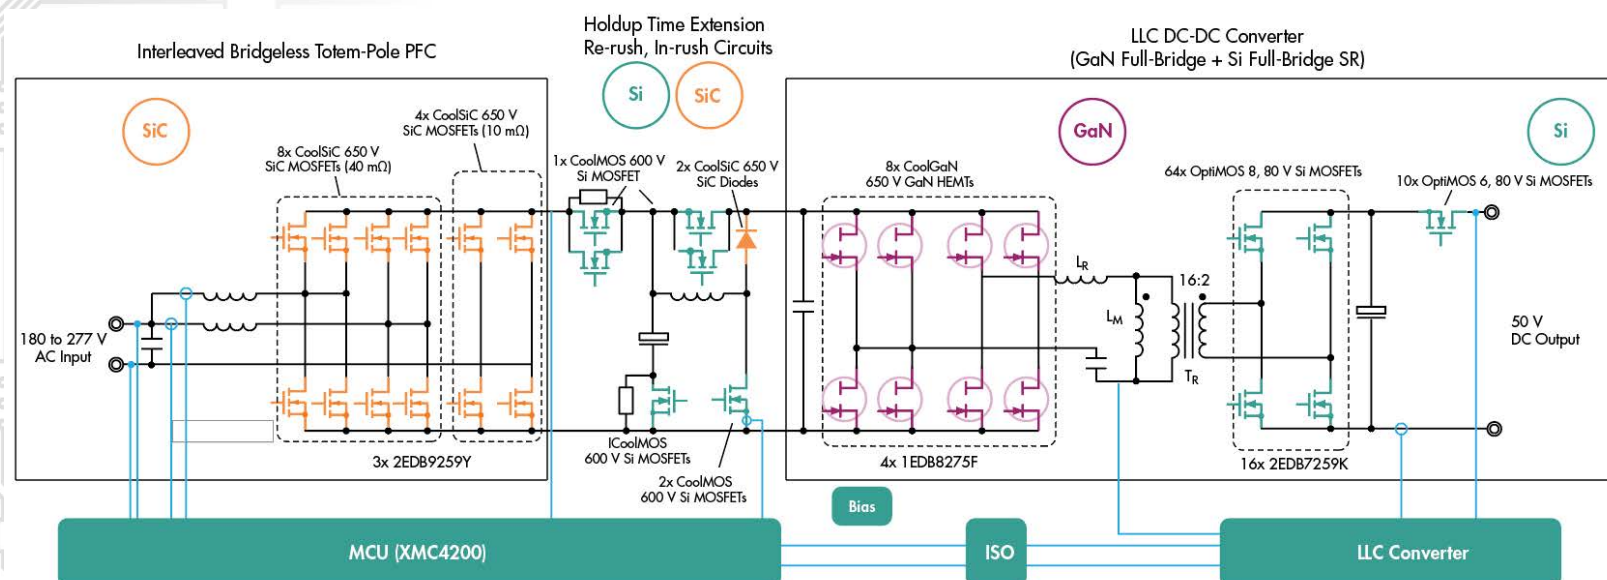
The back-end stage is an isolated LLC converter that regulates the 50-V DC nominal output. The primary side of the isolated DC-DC converter is implemented with eight 650-V GaN power devices to reduce driving losses and switching losses at high frequencies. The bridge is located on a separate daughtercard and drives a pair of high-frequency, series-connected transformers. The synchronous rectifier on the secondary side of the DC-DC converter is based on 80-V OptiMOS silicon MOSFETs.

To meet server hold-up time requirements, the design includes a more advanced baby-boost converter to support a maximum 20-ms hold-up time at 100% of output power while the PSU is losing input AC voltage. During normal operation, the auxiliary boost converter—also referred to as the hold-up time extension circuit—is idle and is bypassed by a 600-V CoolMOS SJ MOSFET.

Overall, the 8-kW reference design achieves a 97.5% peak efficiency and maintains a minimum efficiency of 96.5% across the 30% to 100% load range at 230-V AC input, including fan power consumption. Its outer dimensions of 73.5 × 450 × 40 mm correspond to a power density of 100 W/in³—twice that of the OCP ORv3 specification. Additional efficiency is achieved through digital control, integrated magnetics, and optimized thermal design.

The 12-kW PSU: From Single- to Three-Level Converter Topologies

The next stage in the roadmap is [a single-phase 12-kW PSU reference design](#) that reflects the latest advances in high-efficiency power conversion for data centers. The design will feature a modular architecture composed of a pair of 6-kW modules, each with a 1/2 U form factor, enabling them to be combined for scalable output power. This modu-



3. This is a simplified schematic of the 8-kW PSU reference design, the REF_8KW_HFHD_PSU. Infineon Technologies



larity bolsters light-load efficiency, simplifies manufacturing, and improves power density.

The PSU will adopt a fully three-level conversion topology in both the PFC and the LLC converter stages, enabling operation at higher input voltages by leveraging semiconductor switches with lower voltage ratings.

The two-level totem-pole PFC in the previous generation will be replaced by a three-level flying-cap totem-pole PFC based on 400-V CoolSiC MOSFETs. One of the benefits of a multi-level topology is that it can effectively multiply the switching frequency of the power MOSFETs by the number of levels minus one. As a result, the three-level flying cap topology can reduce the PFC choke size and improve EMI. That not only improves power density, but it also helps thermal performance by enabling more efficient airflow from the internal fan.

Unlike traditional LLC-based topologies, the DC-DC stage operates as a DC transformer, offering faster dynamic response to accommodate rapid load changes from GPUs while maximizing efficiency. These GPUs can require huge amounts of currents when leaping to full power, resulting in large load transients.

To limit voltage overshoot and undershoot, the DC-DC stage must be able to adjust the output voltage as fast as possible. The three-phase LLC converter will use WBG semiconductors to run at faster switching frequencies, increasing the control-loop bandwidth and giving it more dynamic control of the output voltage.

More advanced digital controllers are also used to stay on top of these power spikes, which can result in load transients as fast as 2.5 A per microsecond at the PSU level. In the 12-kW power supply, digital control is managed through Infineon's PSoC 3 and XMC microcontrollers. These MCUs provide real-time monitoring, predictive maintenance, and remote diagnostics—further improving system uptime and operational reliability.

A power pulsation buffer circuit takes the place of the baby-boost converter in the other designs. It smooths out input current from the grid and reduces the bulk capacitance required to meet hold-up time specifications.


The use of advanced thermal management and high-performance components enable the 12-kW PSU to achieve a power density of more than 100 W/in³, with outer dimensions of 69 × 720 × 40 mm.

Transforming PSU Design with Wide-Bandgap Devices

These PSU reference designs represent a clear upgrade path from 3 kW and 3.3 kW to next-generation 8- and 12-kW designs, supporting both greenfield deployments and the replacement of legacy AI power supplies.

But they also reflect a broader industry shift toward digitally controlled, high-efficiency power systems designed for next-gen AI workloads. By strategically combining CoolMOS, CoolSiC, and CoolGaN devices in power topologies that are being reshaped and refined over time, these solutions optimize thermal management and system reliability at scale.

As data centers push for greater density, energy efficiency, and carbon reduction, Infineon's reference designs offer a foundation for high-performance, sustainable power delivery.

to view this article online,  [click here](#)

 **BACK TO TABLE OF CONTENTS**