# Increasing the Power Density and Efficacy of Datacenters Using a Two-Stage Solution for 48V Power Distribution

By Qian Ouyang, Rohan Samsi, and Jinghai Zhou


## Abstract

Current datacenters typically have a 12V backplane and distribution on board and need to convert the voltage down to around 1V, which is usually achieved with a synchronous buck regulator, single- or multi-phase. The racks in these datacenters normally max out at a power rating of 20kW. There is a need in the industry to increase the power density per-rack to about 100kW to reduce the size of these datacenters. This can be achieved by using a 48V backplane and distribution, but there are challenges associated with this approach. To drive 48V down to the board, traditional synchronous buck regulators cannot be relied on. How, then, does one increase the density of datacenters without also increasing cost? This paper outlines a two-stage solution to drive 48V down to the point of load (POL, approximately 1-5V) in a flexible, scalable, and cost-effective manner that will benefit the next generation of server power delivery.


## Proposal

As the demand for datacenters has increased, so has the need for increasing their size and density to accommodate the needs of all users. One of the key limiting factors is the power limitation of servers per rack, which is around 20kW. This limitation arises due to a suboptimal power distribution network. Most mid-planes and backplanes operate at 12V, which requires a large volume of copper and limits the power per rack. The Open Compute Project (OCP) and Google have proposed ideas that would increase the operating voltage to 48V, thereby increasing the installed capacity per-rack to 50-100kW per rack. One of the key reasons why this architecture has not yet found success is due to the lack of downstream solutions (i.e.: 48V down to the board-mounted POL, including processors, memory banks, and other ASICs).

There have been a few different approaches proposed to address the 48V input to POL distribution problem. The main challenges to overcome include scalability, cost, efficiency, and size.

*Scalability and Cost*

It is difficult to distribute 48V to a variety of loads, including small currents used for rails, such as USB and VGA ports, that typically consume a few hundred milliamps each at 2-5V and scale up to the processors, which consume hundreds of amps at close to 1V. Some available solutions include driving the voltage directly from 48V to the load voltages (1-5V) by accurately regulating an intermediate bus and using a DC/DC transformer for the final step-down.

While these solutions are effective for the high-current rails, they are both hard to scale down and more expensive for the majority of the low-current rails, and can even be more expensive for high-current rails as well. Other solutions, such as the use of gallium-nitride (GaN), have been proposed to solve these issues to perform a direct conversion using a simple, synchronous, buck solution. While they do hold significant promise when cost and large volume manufacturability become viable, these solutions at present appear to be distant.

*Efficiency and Size*

The solutions of the board must have both high efficiency and a small size to fit on a current server board. Efficiency of the 48V-to-1V conversion must be at least 93% or above, since the state-of-the-art conversion efficiency is 95% for 12V-to-1V conversion. The size of the 48V-to-1V converters must be no larger than the 12V-to-1V converters due to the limitations of the dimensions of the industry-standard rack and the planar boards that plug into the backplane.

## Solution

The proposed solution for 48V-to-low-voltage distribution is a two-stage process that will improve efficiency, scalability, and cost compared to existing datacenters.

*Stage 1*

The VIN rail (48V) is distributed across the board and subsequently stepped down to a variable intermediate voltage, typically between 5-8V. The variable 5-8V can be generated in clusters for CPU and memory power, while the rest of the power distribution (totaling about 50W) can be generated from a separate converter. The intermediate floating rail ensures complete soft switching, achieving a peak efficiency of 98% by using a half-bridge, resonant, LLC converter. Isolation is not necessary since the input voltage is below 60V. Functional isolation can be achieved more easily by using a transformer in place of an inductor as part of the LLC network. This also aids the step-down from 48V to 5-8V. The fundamental idea is to modularize this first-stage solution (see Figure 1).

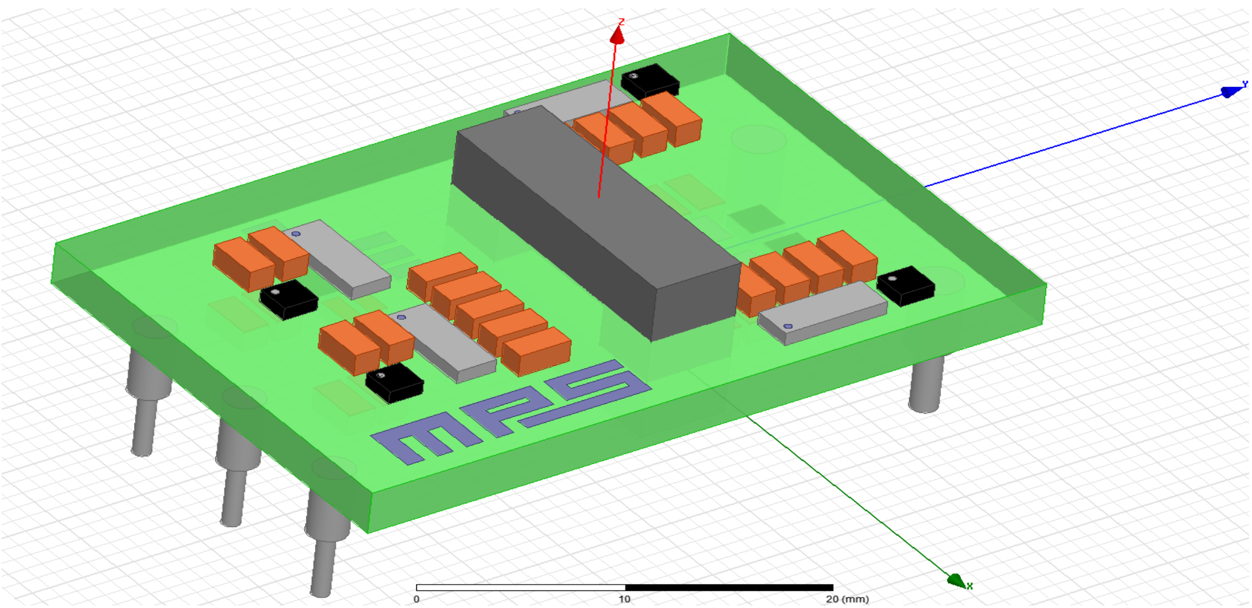


**Figure 1: Front View of First-Stage Module**

The first-stage modules can be scaled as a function of the power delivered, but for a typical single-processor server, only two of these modules are necessary. Another unique feature of this first stage is that it can be multi-sourced. When technologies such as GaN become more prevalent, they can replace these modules seamlessly without affecting the downstream solutions. The variable 5-8V unregulated voltage can also be replaced by a tightly regulated voltage between 5-8V without causing any disruptions to the overall system so that interoperability is maintained.

*Stage 2*

The second stage depends entirely on the power to be delivered. In the case of a milliamp load, the second stage could be as simple as using a linear low-dropout (LDO) regulator. As the power level increases, the second stage could utilize single-phase, synchronous, buck regulators, which are plentiful. With the reduction in input voltage, the need for low-duty ratios is reduced, and the field-effect transistors (FET) and efficiency can be optimized, reducing losses. By reducing the need for high breakdown voltage FETs as compared to the typical 12V rail, the cost of the devices is reduced, and their figure of merit is improved for efficiency. For higher current solutions for processors and memory, interleaved multi-phase regulators can be used (see Figure 2).

With the reduction in input voltage, these multi-phase converters can achieve a peak efficiency of about 97%. With the improvements in feedforward control in most of these converters, the floating input voltage (5-8V, output of the first stage) can be handled easily. The size of these converters is made smaller by using high-frequency conversion, which uses smaller inductors and fewer capacitors.

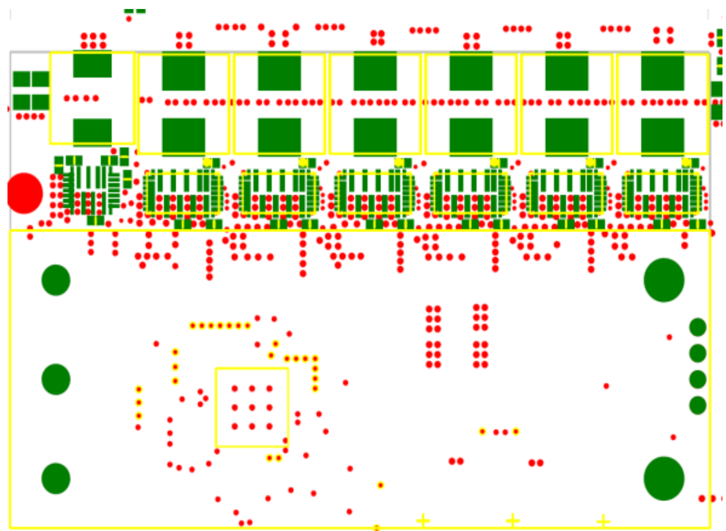

**Figure 2: Second Stage**

## Summary/Conclusion

The overall efficiency of this solution would be about 95%, which exceeds the 93% target for 48V-to-1V conversions and matches the 12V-to-1V state-of-the-art conversion efficiency. The size of the board is not increased since the module can be mounted vertically. The subsequent gains of the size reduction of the second stage account for the increased size of the first stage. The flexibility of using any second-stage converter and correspondingly sizing the first-stage converter increases scalability. With this solution, datacenters can achieve 100kW per rack of density without increasing cost and size.