

Powering your Open Compute ORV3 Style Rack in today's AI-driven landscape



muRata
INNOVATOR IN ELECTRONICS

Powering your Open Compute ORV3 Style Rack in today's AI-driven landscape

The Open Compute Project (OCP), initiated in 2011 has indeed revolutionized the way we think about data center hardware. By applying open-source principles to hardware design, OCP has fostered collaboration and innovation across the industry. By making designs available to everyone, OCP has significantly reduced costs for data center operators, allowing them to build more efficient infrastructures.

Over the past decade, OCP has become a leader in driving innovation across various hardware domains in rack power and backup as well as network equipment, all forms of IT gear, CPUs, GPUs, and storage devices.

The transition from a chaotic setup with multiple power supplies in an N+N configuration to the more streamlined N+1 configuration in OCP racks has indeed transformed power management. This shift not only simplifies the physical setup but also contributes to better overall performance and reliability in data centers.

Before OCP, typical racks used two power supplies per server with AC cables running everywhere — a setup known as N+N redundancy. In an N+N configuration, two power supplies are installed for every single supply required. This ensures full redundancy: if one fails, the other instantly takes over to keep the server running without interruption. Under normal conditions, each power supply carries 50% of the load, balancing power delivery and extending component life.



The initial OCP Style rack had a single busbar and two AC cords per single power shelf in an N+1 configuration.



The typical OCP rack power supplies are six-across design in a 21" shelf maximizes space utilization, allowing for more power supplies in a compact footprint. In a Murata OCP style 21" 10U power shelf, with two 3-phase inputs powering three power supplies each, it eliminates the complexity of balancing loads across multiple phases, which was a challenge in the previous N+N configurations. Each power supply drawing from a single phase of the 3-phase input enhances reliability, as it reduces the risk of overloading any single phase.

In the delta input configuration, power Supplies 1, 2, and 3 draw from one 3-phase AC input, while supplies 4, 5, and 6 draw from another. Each power supply operates line-to-line, which can enhance efficiency and reduce the risk of neutral currents. In a Wye configuration, each power supply connects to a different single phase to neutral. With all six power supplies sharing the same output voltage and current, the system achieves excellent load balancing across the phases. That means all six power supplies have the same input current making them all balance the 3-phase input currents.

The balancing of input currents is just one of many advantages of the OCP rack power design.

In an N+1 configuration within a typical OCP power shelf, each power supply contributes 1/6 of the total output power to the busbar during normal operation. This design ensures that if one power supply fails, the remaining five automatically share the load, maintaining power without interruption.

In the ORV2 rack, the single busbar was typically used to distribute 12V to the entire rack. Murata has a few different versions of the ORV2 power shelf, with six power supplies, each drawing from three-phase inputs. The ORV2 has a solid busbar running down the back of the rack ensures efficient power distribution to all servers, minimizing resistance and potential power loss. The requirement for precise alignment and secure connections (using nuts and bolts) between the power shelf and the busbar in ORV2 is essential for handling high currents, especially with up to 15kW of power translating to 1250A.

The biggest issue with 12V and high currents in the OCP rack busbar is the power loss which also translates to thermal rise in the busbar and connections. The power loss, for a given

section of busbar, is equal to I^2R where I is the current in the section of busbar and the R is the resistance in that section of the busbar.

Due to this reason, Google transitioned to 48V/ 54V systems in their OCP style racks. Google shared the 48V rack power architecture design to the OCP community in 2018. They learned long ago that by moving to 48V in a rack, the I^2R power loss was reduced dramatically. To understand this, we first look at the difference between the current in an equivalent power system when you move from 12V to 48V. By increasing the voltage from 12V to 48V, the current required for the same power output decreases by a factor of four. Since the power loss is proportional to (I^2R), moving to 48V results in a drop that is only 1/16th of what it would be at 12V. This significantly reduces both the power loss and thus the thermal rise in the busbar and connectors.

The use of a DC-to-DC converter to step down the 48V input to 12V for server components is a crucial aspect of modern power architecture. This setup allows servers to efficiently utilize the higher voltage while still powering sensitive components like processors that require lower voltages. Let's stop here and look at the 48V.

First, it's not usually 48V that is used. It's 54V that Google has been using for years. 54V was common in telecom applications, especially for backup power systems. The transition from 48V to 54V makes sense given the charging characteristics of lead-acid batteries. The 2V lead acid cells were stacked to equal 48V. Then when the lead acid cells were fully charged, the stack equaled 54V. Central office power supplies usually allow a nominal 54V input; typically, 40V to 70Vin. That is where the 54V stems from.

If the servers DC input can handle a 54V input, everything works fine. It can be confusing when people refer to it as a 48V system architecture.

Meta decided to implement a 51V system in their Open Rack V3 (ORV3) architecture, rather than sticking strictly to the traditional 48V. What could be the reason behind this step?

Telecom 48V = 54V

Google 48V = 54V

Meta ORV3 48V = 51V/ 50V

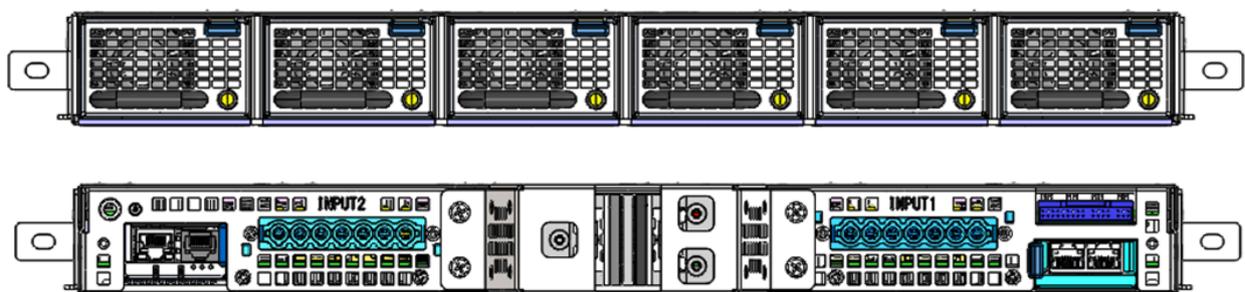
Running most racks at less than 8kW while having a maximum capacity of 13kW provides flexibility and efficiency, especially in managing peak loads. By using a non-isolated 4:1 DC-to-DC converter, Meta can efficiently step down the 51V to 12V while minimizing energy losses.

This design choice is particularly beneficial for Meta to optimize energy consumption across their vast server farms.

Today the ORV3 standard power shelf is 6 across 3kW power supplies for 15kW +1, or a total of 18kW per shelf at 51Vout. This is referred to as N+1. This means the N is 15kW with an extra +1 or a 3kW power supply for redundancy if one of the power supplies fails. The system can continue to operate producing up to 15kW until the failed power supply is hot swapped/ replaced.

Murata's present ORV3 shelf is a 21 inch, 6 across, 3.6kW power supplies. This is 18kW +1 for a total 21.6kW per shelf at 51V or 54V. The flexibility of using either 51V or 54V in the Murata ORV3 shelf is a smart design choice, catering to different server requirements while maintaining high efficiency. Murata built our 3.6kW, 54V power supply, years ago in a compact 68mm width to allow for 6 across in a 19-inch shelf. When we optimized the power supply for the ORV3 51V, we retained the 3.6kW in the 68mm wide package that allows for 6 across in a 19 inch or 21-inch shelf. The optimization for 51V required design changes to the magnetics and the control loop to keep the efficiencies flat across the power range as required in OCP products.

If the customer wants dual input redundancy at the power shelf level, which is not a part of the typical OCP architecture, it can be accomplished with a different AC input for each 3 power supply groups. This means power supply 1, 2, and 3 will have AC input A supplying the power and power supplies 4, 5, and 6 will have AC input B supplying power. All the outputs will be tied together and shared. When AC input A fails and goes away then power supplies 4, 5, and 6 will keep operating. There is no switching involved like a typical ATS as the power supplies 4, 5, and 6 will simply pick up the required output load and keep operating seamlessly. This seamless transition ensures that ORV3 can have redundant AC inputs assuming the output power of the 3 power supplies is enough to handle the load. In the case of the Murata ORV3 shelf, this implies 10.8kW per 3 power supply group.



If more than 10.8kW is required, then you could use 2 shelves with AC input A powering 6 of the power supplies and AC input B powering the other 6 power supplies. The shelves would operate in the same manner together giving you 21.6kW of power with redundant AC.

In today's world of AI Architectures using a typical Nvidia GH200 GPU, a rack full of these devices would require 60kW of power, 3 or 4 of the present Murata power shelves. Given that the shelves are presently designed to allow up to 3 shelves to run in parallel, it is best to split the rack into 2 power zones with 2 shelves running in parallel in each zone for a total of 39.6kW +1. Two Murata's power shelves would provide 39.6kW versus 33.0kW per power zone with most ORV3 present power shelves.

To address the growing power demands of AI-optimized racks, Murata and its partners have introduced the ORV3 HPR 5.5 kW power supply, optimized to fit six units in a 21" ORV3 HPR shelf. With the fully populated 33 kW shelf now in production, data centers can reliably deliver the 120 kW required to power a full rack of Nvidia GB200 GPUs.



Grace Hopper Data Center Racks: "Data Center Frontier"

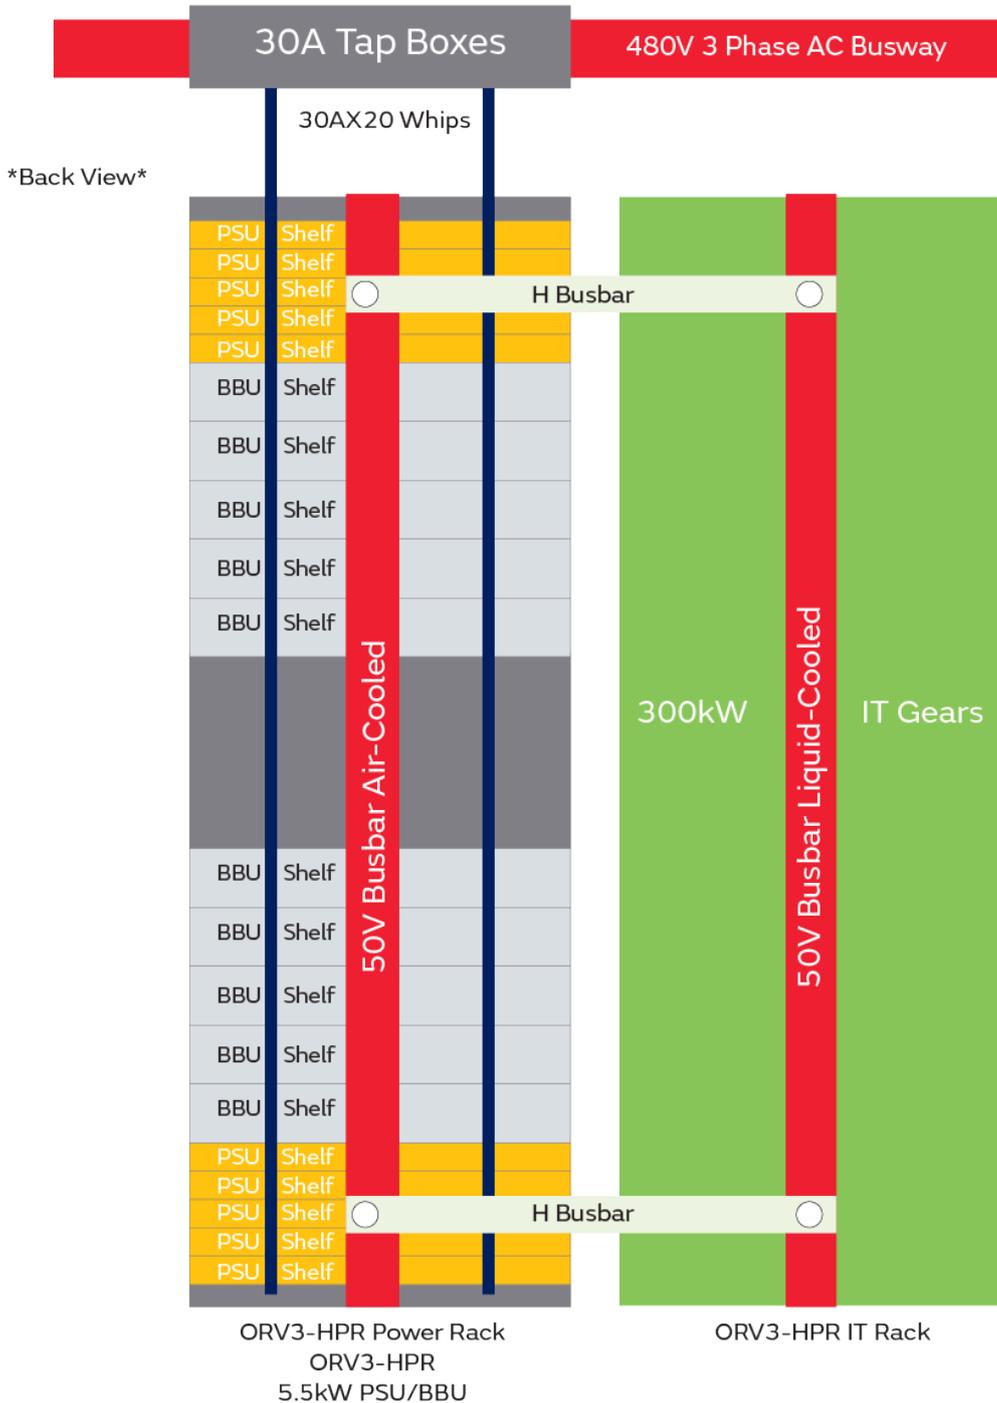
The shift towards water cooling using a Direct to Chip (DTC) cold plate design for high-power racks represents a transformative change in data center cooling approaches. This innovative cooling method has fundamentally altered the architecture of cooling systems within data centers. While this transition to water cooling is significant, it is essential to note that the power supplies will continue to utilize air cooling. Given that they only dissipate about 4% to 5% of total power, air cooling remains a practical and efficient solution for these components.

To meet the extreme power requirements of AI-optimized, GPU-dense racks, the HPR ORV3 system increases both the size and conductive efficiency of its busbar, enabling up to 93 kW of reliable power delivery. This enhancement, when paired with robust thermal management, safeguards operational stability and maximizes energy efficiency under sustained high loads.

As computing power demands surge, Nvidia projects that data center rack capacities will double every 6 to 12 months, potentially exceeding 280kW. To address this, the Open Compute Project (OCP) has developed a compact 12kW power supply with a 50V output for the ORV3 HPR2 21" 10U shelf. With an innovative design measuring 83.4mm wide and a slim 15mm side-mounted controller, this setup delivers up to 72kW per shelf, indicating that future AI workload may require four or more shelves to meet growing energy demands

Due to the 280kW power limitations, integrating all GPU servers, power shelves, and battery back-up units (BBUs) into a single rack has become impractical. As a solution, the decision has been made to separate the GPU servers into dedicated Compute Racks while relocating power shelves, BBUs, and capacitor banks into dedicated Power Racks. This innovative configuration, known as the Mount Diablo architecture, is being collaboratively developed and implemented by industry leaders such as Google, Microsoft, and Meta. These organizations are pooling their expertise to advance and optimize future designs within this framework.

The Mount Diablo 50V ORV3 HPR3 power architecture employs solid DC busbars to efficiently transfer 50V between Power and Compute Racks. In GPU-intensive environments, simultaneous accelerator startups can create inrush currents that exceed continuous PSU capacity. To manage these transients, the system integrates a high-capacitance energy storage module—utilizing supercapacitors or advanced supercells—that delivers rapid surge currents without impacting upstream AC lines. This "AC input smoothing" or "peak shaving" approach stabilizes voltage, preserves PSU performance, and mitigates grid disturbances, ensuring reliable operation for next-generation, high-density compute loads.



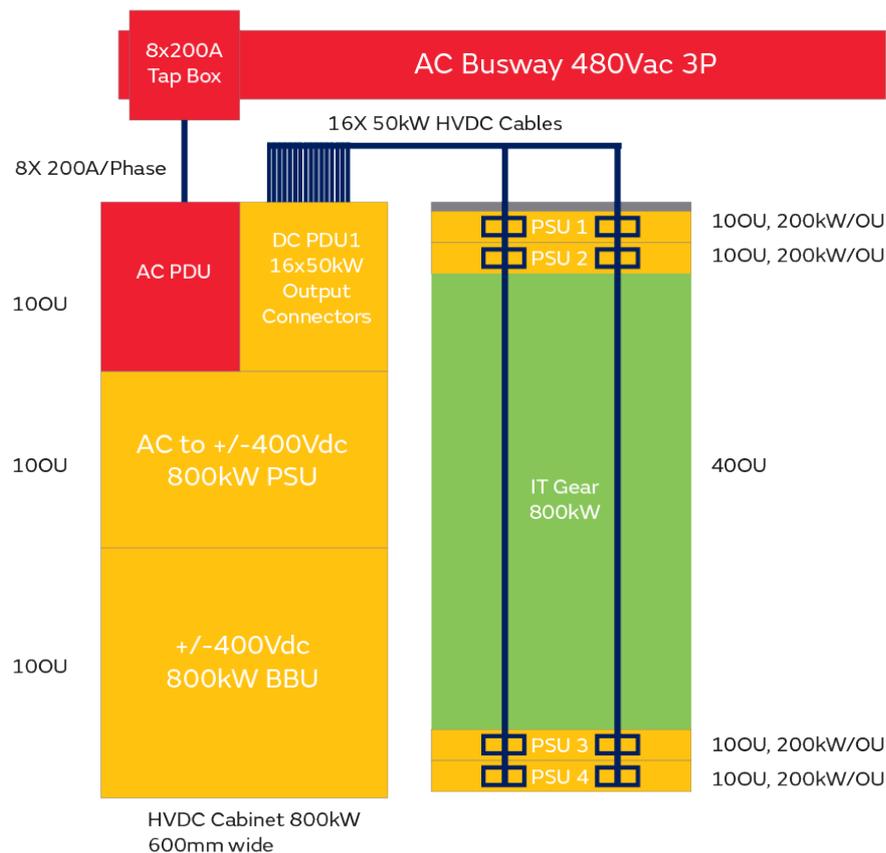
Current single-phase input topologies with 48V/50V/54V configurations are approaching efficiency limits. To address this, transitioning to a 3-phase input design for individual power modules is the next logical step. Additionally, increasing the output voltage to 800VDC will help reduce current levels in output connections and cables, enhancing overall system efficiency.

The Case for 800VDC: A Standard in Vehicle Power Infrastructure

The automotive industry is increasingly adopting 800VDC as a standard for their power systems, including batteries. This shift is supported by a robust ecosystem of connectors and components designed for this voltage configuration, enhancing performance, efficiency, and compatibility across vehicle platforms.

Due to the current semiconductor supply limitations, OCP is shifting to a +/-400V configuration with a center reference point, instead of transitioning to 800VDC. This change reduces the current by a factor of 8, leading to a significant decrease in voltage drop and temperature rise—consistent with the I^2R relationship.

Transitioning to a 400VDC voltage bus introduces new challenges that require careful assessment, particularly concerning safety. Higher power levels lead to increased input current, which must be managed effectively. This shift in output voltage is essential to support the anticipated 600+kW demands of future GPUs, such as those announced by Nvidia for the Vera Rubin Ultra configuration.



GPU servers will continue to operate at 50 V input, requiring 400 V to 50 V DC-DC converters or a dedicated converter shelf within the compute rack. This pivotal shift is opening exciting opportunities at the intersection of advanced power systems and AI.