

# Application Note



## Applying Open Compute Rack and Power Specifications to New and Existing Data Centers

[belfuse.com/power-solutions](http://belfuse.com/power-solutions)



**bel** POWER  
SOLUTIONS &  
PROTECTION  
a bel group

# APPLYING OPEN COMPUTE RACK AND POWER SPECIFICATIONS TO NEW AND EXISTING DATA CENTERS

When Facebook and others launched the Open Compute Project (OCP) in 2011, the goal was to accelerate the commercialization of hardware designs that enable greater data center efficiency and flexibility. Since then, the OCP has gained traction in the market, growing into an influential community focused on open data center hardware.

Among the major developments in the last several years are the addition of notable members, including Apple, Cisco, Nokia and Google; continued growth in the number and maturity of OCP specifications; and the development of specifications for telecommunication data centers.

From its beginning, the OCP has questioned and reconsidered traditional approaches to data center design, challenging the idea that we should continue doing things because that is the way they have always been done. That's great if you're just building out your computing infrastructure, as several of the key players in OCP are, but it can create challenges for organizations with significant investments in legacy infrastructure. Fortunately, as the community has grown, OCP specifications have expanded to increase their applicability in legacy environments.

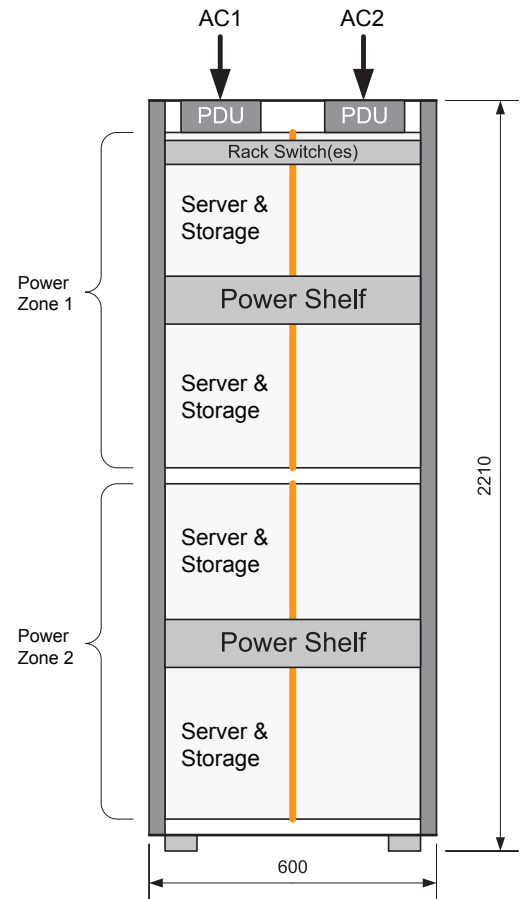
This is causing data center operators across the spectrum to take notice. They are increasingly evaluating the viability of OCP specifications when building new data centers or expanding existing facilities. They often find opportunities to drive down costs and improve efficiency, but also face some key decisions on how best to apply the specifications—or the principles behind the specifications—to their particular environment and objectives.

## The Changing Shape of Equipment Racks

One of the shifts the OCP made from conventional data center standards in its initial Rack and Power specification was to redefine the dimensions of the IT equipment rack from the traditional 19-inch to 21-inch wide racks.

The 19-inch rack, according to the OCP, didn't provide enough space to accommodate the airflow and serviceability requirements of today's IT equipment. Adding two inches to the width of the rack creates more room for cable runs, simplifies server deployment and service, and improves airflow to create opportunities for more efficient data center cooling. Facebook has standardized on the 21-inch rack in their data centers and this has led to the commercialization of racks and supporting products that comply with the 21-inch specification.

The OCP Rack and Power specification, of course, went beyond rack dimensions. The specification essentially unbundles the traditional server configuration, deploying servers, storage and power on separate trays within the rack. With these components physically separated, servers can be upgraded more cost-effectively when new processors are introduced. It also provides greater flexibility in configuring racks with required equipment.



OCP V2 rack showing location of servers, storage and power shelves.

# APPLYING OPEN COMPUTE RACK AND POWER SPECIFICATIONS TO NEW AND EXISTING DATA CENTERS

These benefits were compelling enough to cause OCP members who weren't ready to abandon their 19-inch racks to apply the principles of the Rack and Power specification to EIA-compliant 19-inch racks. Now, designers and operators seeking to leverage the benefits of the OCP specification can choose between 19-inch and 21-inch racks, with a full complement of racks, servers, storage and power products available for each.

## Accommodating Higher Density

The OCP specifications were driven by the needs of large-scale deployments, mostly greenfield data centers built in areas where real estate costs are low. In these environments, capacity can be easily spread out and there isn't a good economic case for high-density racks.

As a result, the OCP specification supports 13.2 kW of IT equipment within the rack. This is accomplished through the use of two power shelves within the rack, each of which consists of three power supplies in an N + 1 configuration. Each shelf supports a different zone within the rack and is connected to the rack management controller or facility-level monitoring to collect data on temperature, power in, power out, and utilization. Power shelves can also include "smart boxes" which archive data in the event of a failure for diagnostics.

While there generally isn't room within the rack for additional power shelves, you can accommodate higher densities within the OCP rack by choosing a power shelf designed for high density. Shelves are available that include six 3 kW power supplies in an N + 1 configuration that uses just 1U of rack space. This effectively doubles the power capacity of an OCP rack.

Due to the efficiency curve of modern power supplies, these higher density shelves can be deployed without compromising efficiency, even if the initial capacity of the rack is well below the design specification. The power supplies have a maximally flat efficiency curve across the load range.

The efficiency curves also highlight the differences between Platinum and Titanium power supplies. Data centers striving to achieve the highest levels of efficiency will be tempted to specify Titanium power supplies. However, it's

useful to evaluate the efficiency curves in light of your expected requirements before making a final decision. While Titanium power supplies provide efficiency improvements at certain points of the curve, the efficiency of Platinum power supplies may provide improvements at other points, such as light and full loading. Comparing the performance of each at the points on the curve you plan to operate can help guide decisions around whether the additional cost of a Titanium power supply is warranted by the efficiency gains that will be realized.



Efficiency curves for the Bel Power Platinum (PFE3000) and Titanium (TET3000) power supplies.

# APPLYING OPEN COMPUTE RACK AND POWER SPECIFICATIONS TO NEW AND EXISTING DATA CENTERS

## The Question of Voltage

One of the challenges designers now face in configuring data center power is making decisions about input voltage to the rack and output voltage from the power shelf to the servers. In the past, 208 V to the rack and 12 V to the server were common, but as designers seek to minimize losses within the critical power path created by distribution, conversions and step-downs, alternate approaches are being considered.

The OCP specification addresses this by bringing 277/480 V power (230/400 V in Europe) directly from the transformer to the rack, eliminating the stepdown to 208 V. This can create a 2.5 percent efficiency improvement compared to 208 V without making significant changes to the way power is delivered to the rack.

Some designers see an opportunity to go even further. Why not, they ask, simply rectify the 480 V power coming into the data center and deliver 380 V DC to the rack? This saves an additional 1 to 2 percent compared to the OCP specification, but there is resistance among operators to running high-voltage DC in the data center. In addition, equipment such as fans may still require AC power, creating complications. Still, the efficiency of DC power is appealing and it would not be surprising if this approach finds more proponents in the future. The power shelves used in an OCP rack can be configured to accommodate 208 V AC, 277 V AC or 380 V DC.

Output voltages can also vary. The initial OCP specification provided for 12 V output from the power shelf, which is what most server hardware is designed to accept. But when Google joined the OCP in 2015, they developed an OCP specification for 48 V output from the power shelf. Google is a proponent of 48 V servers for the same reason the OCP opted to deliver higher voltages to the rack—the higher the voltage, the greater the efficiency.

In this case, delivering 48 V to the server reduces the voltage drop across the busbar and connectors by a factor of 16 (assuming same conductor cross section), increasing the overall efficiency compared to 12 V. The bigger advantage of 48 V, however, is that it provides a path to the future for racks with much higher power requirements. This cannot be easily realized with 12 V distribution because of the higher currents required.

Until Google joined the OCP, the only servers that could accept 48 V input were those manufactured for Google. It's still too early to determine whether the OCP 48 V specification will drive broad commercialization of 48 V servers, but the initial signs are promising. If you're seeking to achieve the highest possible efficiency and are willing to live with the limits of the current market for 48 V servers, this could be a viable option that is supported by the current generation of power shelves.

A further variation comes into play as telecommunication companies transition to data centers. Telco providers have a long history with -48 V input equipment. As they shift from central offices with purpose-built hardware to data centers with standardized hardware and software-defined services, they will increasingly replace their specialized equipment with servers running at 12 V with an intermediate -48 V to 12 V DC/DC shelves. In future, they may migrate to servers directly supplied with 48 V.

# APPLYING OPEN COMPUTE RACK AND POWER SPECIFICATIONS TO NEW AND EXISTING DATA CENTERS

## Making the Move to OCP

As the OCP specifications continue to mature, they are being applied to a wider range of environments and requirements. With a mission of improving the efficiency and flexibility of data center hardware, the OCP is helping advance new approaches to data center design that have broad benefits within the industry. While not everybody is well positioned to adopt the OCP specifications, data center designers and developers should, at minimum, stay abreast of the developments within the OCP and how they are shaping the future of data center hardware.

**Bel Power Solutions**, a Bel group, offers world-class AC-DC and DC-DC power conversion products, value-add solutions, complete box-build solutions and contract manufacturing services. Applications of our power conversion devices range from board-mount power to system-level architectures for servers, storage, networking, industrial and telecommunications industries.

Bel is headquartered in Jersey City, New Jersey. Founded in 1949, Bel designs, manufactures and markets products that power, protect and connect electronic circuits.

[belfuse.com/power-solutions](https://belfuse.com/power-solutions)



**bel** POWER  
SOLUTIONS &  
PROTECTION

a bel group

[belfuse.com/power-solutions](http://belfuse.com/power-solutions)

Asia-Pacific  
+86 755 298 85888

Europe, Middle East  
+353 61 225 977

North America  
+1 408 785 5200

© 2018 Bel Power Solutions & Protection

App Note BCA.00257 Rev. AA