Open Compute Project

Open Rack Hardware v0.6

Authors:
Pierluigi Sarti, System Designer
Jon Ehlen, Mechanical Engineer

1 Scope

This document describes the technical specifications for the custom rack that houses Open Compute Project server technologies. It includes an example of Facebook's Open Rack implementation under the Open Compute Project — an Open Rack triplet with 9 power zones.

2 Contents

3 Overview

When data center design and hardware design move in concert, they can improve efficiency and reduce power consumption. To this end, the Open Compute Project is a set of technologies that reduces energy consumption and cost, increases reliability and choice in the marketplace, and simplifies operations and maintenance. One key objective is openness—the project is starting with the opening of the specifications and mechanical designs for the major components of a data center, and the efficiency results achieved at facilities using Open Compute technologies.

The main components of this project are the Open Rack and equipment chassis that can be configured as a server rack, storage box, and more. The Open Rack uses an all-encompassing design to accommodate compatible Open Compute Project chassis components, and include the power solution as well as input and output voltage distribution.

3.1 License

4 Open Rack Overview

This Open Rack implementation is a triplet configuration — a three column rack in a single cabinet.

Each column in the rack is divided into one or more *power zone*s (3 power zones for this implementation). A power zone comprises an *equipment bay* for the compute, storage, or other components and a *power shelf*, which powers the compute components in the equipment bay. The power shelf includes backup power capability using an external Open Compute Project Battery Cabinet, or it comprises a power shelf and a Battery Backup Unit that embeds batteries. If there are multiple power zones, they are stacked one above another. The power shelf is described in Power Shelf Specifications.

Each column has space for up to three Ethernet switches. In this implementation, the switches are installed in one of two locations:

- All three switches are at the top of the rack, above the topmost power zone.
- Each power zone has its own switch located within its equipment bay, typically above the power shelf, so the switch may be powered at 12V from its own power zone.

Users may arrange the switches and equipment in other configurations as needed.

Each equipment bay in a power zone can be set up in various configurations and can accommodate different numbers and height of equipment chassis (such as individual server or storage chassis). Motherboards and other Open Compute hardware are hot-swappable within the individual equipment bay, as they take 12V power from the rack.

Variable configurations are achieved because the rack enclosure is designed to accommodate various equipment mounting configurations. The rack has no external side panels; rather, shelves in the equipment bay are locked in place to the vertical side walls and completely block the side access and air. As long as the design follows Open Compute Project standards, the chassis — whether server chassis, storage chassis, or other — will fit in the rack slots and will be powered.

4.1 Open Rack Triplet Images and Drawings

Figure 1 Open Rack Triplet

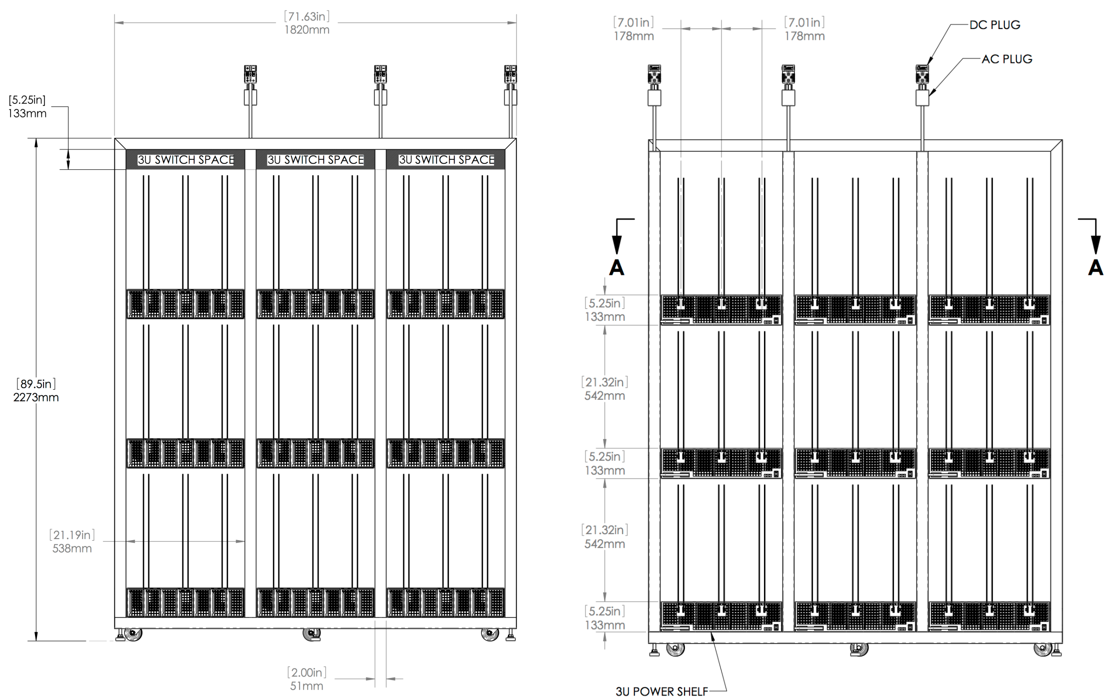

Figure 2 Open Rack Structural Drawing with Vertical PDUs

4.2 Mechanical

The height of a slot in an Open Rack is measured in *OpenU*. Unlike a *rack unit* (RU), an OpenU is 48mm high, which includes the space between each chassis unit. The chassis unit height is measured in 0.5 OpenU increments, starting from 1xOpenU, for flexibility in form factor. In this implementation, a single chassis can be between 1xOpenU and up to 12xOpenU high (maximum), in 0.5 OpenU increments. Or there can be multiple chassis, as long as all the chassis add up to 12xOpenU per power zone.

The side walls of each column provide supports for the chassis using simple L-shaped metal bent, directly on the walls, and customized to the chassis heights; the walls are easily swappable with new ones.

The Open Rack triplet has these dimensions:

- The rack is nominally 1820mm (±2mm) wide, which includes up to three 538mm (+1/-0 mm) wide equipment bays per column.
- Depth can be no shorter than 914mm and no longer than 1220mm with a maximum depth for IT equipment of 914mm.
- Total height is 2413mm.

4.2.1 Chassis Dimension Guidelines

- Width is 537mm (±0.5mm).
- Height must be a minimum of 1xOpenU, and can grow in height by ½xOpenU (24mm) increments. The chassis height should allow for 2.5mm to provide enough clearance between the slots. For example, a 3xOpenU chassis would be 141.5mm in height: (48mm * 3xOpenU) -2.5mm = 141.5mm.
- Maximum depth is 280mm shorter than the chosen total rack depth.

The power shelf must comply with OCP standards for input and output voltage and mechanical mounting scheme in the Open Rack. The power shelf follows the same mounting rules of a compute chassis, and power shelves can be installed anywhere in the column. In this implementation there are three power shelves, each one installed at the base of its power zone. See section 5 for more information.

The rack has a minimum of six casters. They are mounted in such a manner so they withstand tipping under normal deployment or relocation activity. Six leveling feet are threaded into nuts welded into the bottom of the triplet. PDUs are installed in the rear of the rack, one PDU each column; they distribute AC only, or AC and DC in the same enclosure (DC voltage would be used for backup). Management of IT cables takes place from the front, between the vertical rack posts.

Refer to the CAD model for exact dimensions.

4.3 Rack Height Options

As already mentioned, a single chassis can be between 1xOpenU and 12xOpenU in height, as long as all of the chassis add up to 12xOpenU (maximum) per power zone.

For example, in a configuration with three power zones per column, each power zone can have six 2xOpenU chassis; or one 4xOpenU chassis, three 2xOpenU chassis and two 1xOpenU chassis; or four 1.5xOpenU chassis and two 3xOpenU chassis. Each bay can be partially stuffed as well. If the power solution is single phase — as in this implementation — then the three power zones must be populated identically for phase balance

4.3.1 Suggested Rack Heights

The suggested maximum height of the rack including power and IT equipments is 48xOpenU, where each power zone is 12xOpenU high plus a 3xOpenU power shelf, with 3xOpenU for switches.

The same rack can be 45xOpenU high, where each power zone is 11xOpenU.

The same rack can be 42xOpenU high, where each power zone is 10xOpenU.

There is no maximum height for an Open Rack; it can be any OpenU in height. It can have any number of power zones per column and also any number of OpenU allocated for the rack switches. Basically, the power zone defines the Open Rack standard, and not the rack itself.

4.4 Prescribed Materials

The equipment chassis is zinc-plated sheet metal. The chassis allows for the easy installation of hardware components without requiring any tools.

Primary materials used in construction of the rack include the following:

- Cold-rolled steel in sheet form
- Zinc pre-plated cold-rolled steel in sheet form
- Toolless sheet metal without any threaded holes
- Plastic cabling ducts

All joining in the rack is done by welding. The raw steel square tubing is welded together to form the frame. After welding, black powder coat paint is applied to the weldments. Typical square tubing is 50.8mm x 50.8mm.

Next, zinc-plated steel side walls are screwed onto the rack sides using threaded fasteners. The side walls are installed on both outer sides and in between each column. The side walls can be customized for the chassis height and are easily swappable.

One panel is installed at the top of each column; one rivet in each corner secures the panel to the top of the rack. A zinc-plated panel is installed at the bottom of the rack and is held in place by rivets or screws. The top and bottom panels are baffles that close off air holes for airflow impedance.

The cabling duct is mounted to the vertical columns as needed. It allows cables to enter and exit the rack.

The rack is shipped to an assembly facility where the electrical components are installed and network switches put into place. Server chassis (or other equipment such as storage chassis) are installed at this time.

4.5 Thermal Specifications

The following table indicates the thermal specifications for the rack.

Thermal Specification	Value
Loading	Idle to 100%
Inlet temperature	18°C to 35°C
Humidity*	Approximately 30 - 90%
Altitude	1000m (~ 3300ft)
Dew point*	5.5°C minimum

Figure 3 Thermal Specifications at Rack Level

* Based on regional climate conditions where the data center operates; supply air temperature 18°C to 27°C, relative humidity 65% maximum, dew point 5.5°C to 15°C). For

more information, including a psychrometric chart, see the *Open Compute Project Data Center v1.0* specification.

5 Power Shelf Specifications

The Open Rack can be powered either by three-phase (3 or 4 wires) or single-phase AC voltage, or high voltage DC (HVDC), depending upon the power solution adopted. Typical three-phase voltage is 320Vac-530Vac. Typical single-phase voltage is 180Vac-305Vac. Output voltage is 12Vdc nominal — but can be slightly higher to compensate for some distribution voltage losses — and is connected directly to the bus bars. The positive and negative output voltage terminations are electrically isolated with the rack chassis ground.

The power solution must support redundancy, no matter the topology or input voltage adopted, and regardless of how much space it occupies in the power zone. A battery backup unit (or any backup scheme used to support the system in case of an AC outage) cannot be used as a redundant power scheme.

Dimensions (which are the same whether or not a BBU is integrated with the rack):

- Width: 533mm ±1mm. This width is slightly less than a chassis to allow possible installation of the power shelf with fixed L brackets pre-installed on rack posts.
- Height: Up to 3xOpenU minus 2.5mm.
- Depth: Maximum depth should be at least 150mm shorter than the maximum compute chassis depth, to provide enough room for the 12V connection from the power shelf to the bus bars.

5.1 Powering the IT Compute Chassis

Three pairs of bus bars are installed at the rear of each of the three power zones in the column. Each pair is connected to the output voltage positive and negative terminations provided by the power shelf installed in the same power zone (see the mechanical drawing for details). The compute chassis plug directly into the bus bars (up to three chassis in the same rack slot) using a blind-mating hot-pluggable bus bar connector clip assembly. Chassis are hot swappable from the front of the rack and are provided with a latch system that locks them in place once fully inserted. The connector clips are rated 80A after 20% derating, are nickel plated, and ensure a contact resistance less than 1mOhm when mated with the bus bar. The bus bars run the whole height of the power zone, in effect providing the ability for a chassis to plug in to the bus bar at any height.

Each bus bar pair is mechanically protected by a cage, perforated for air circulation, for safe operation (no perforated panel is installed in the rear of the rack). The nominal voltage available at the Open Rack bus bars is 12V (suggested voltage is 12.5V). Each chassis includes the clip assembly installed to the rear panel, with positive and negative 12V cables routed to the system boards sitting inside the chassis. Each chassis includes a hot swappable controller and a ferrite choke clamping the 12V cables for high frequency CM noise reduction.

The positive and negative bus bars on the Open Rack are nickel plated, 3mm thick and 17mm apart. They are offset by 2.5mm, allowing a pre-mating of the 12V negative connector chassis clip. From the perspective of the front of the rack, the negative bus bar is 2.5mm closer to the front of the rack than the positive bus bar. The clip can engage with the bus bars up to a maximum of 22mm. The bus bars have a knife shape on the front side to easily allow the mechanical mating clip. The clip assembly is composed of two clips and one plate where the two clips are installed. The plate assembly is installed in the rear panel of the chassis, and can float +/-3mm horizontally and +/-4mm vertically. The plate includes self-guiding fixtures allowing the clips to blind mate with the bus bar, with no extra guiding pins needed at chassis level. This prevents the bus bar plating from wearing out during transportation, when a rack is fully populated with compute chassis before shipping, and compensates for all unavoidable mechanical tolerances for a flexible and effective design.

The 12V output of each power zone is independent and floating with respect to the other power zones; this helps to reduce potential DC currents from looping through the rack, and against onset noise. The positive and negative bus bars are electrically isolated with the rack chassis ground. Each slot can power up to three independent chassis (each 175mm wide), two independent chassis (each 265mm wide), or one chassis (537mm wide). When only one 537mm wide chassis is used, it takes power from the central bus bar. When two or three chassis are used, they take power from the corresponding bus bar pairs; in which case, a 537mm U-shaped shelf is pre-installed in the rack slot to make possible the installation of the individual chassis.

The two or three 12V bus bars available at each slot level in the rack cannot be paralleled inside the same chassis because this would cause unwanted recirculation of DC current between the chassis installed in the same power zone, lowering efficiency and causing potential nuisances especially when chassis are hot swapped. Each bus bar pair can support up to 5kW DC (12V, 400A), which could be increased by adding more copper. This is equivalent to 15kW (3x5kW) per power zone, which is equivalent to 45kW per column, and equivalent to 135kW maximum DC power per triplet. In this example Open Rack implementation the power density is much lower than these maximum levels because this helps the electrical and thermal efficiency at system level. This Open Rack has three power zones of 4.2kW DC each, equivalent to 12.6kW (3x4.2kW) per column, and equivalent to 37.8kW maximum DC power per triplet.

The bus bars, PDUs and power solution can be adjusted for higher power levels, depending on required power per column, efficiency target and cooling conditions. The bus bars, PDUs and power shelves are interchangeable in the rack.

An Open Rack has three bus bar pairs installed in each power zone, but it is also configurable with two or one bus bar pairs in each power zone:

- Three pairs (+/-) of bus bars: on the left, right, and middle of the column
- Two pairs (+/-) of bus bars: on the left and right sides of the column
- One pair (+/-) of bus bars: down the center of the column

A compute chassis designed under OCP should always clear the keep out zone where the three bus bars are installed in the Open Rack when the chassis is fully inserted in the slot. A fully compatible Open Rack always includes all the three bus bar pairs. See the mechanical drawings for more details.

The bus bar system is safety certified to ensure safe operations, and for compliance to the safety standards of the clip mating with the bus bars (at full load vs. maximum temperature rise, maximum voltage drop, and so forth). The bus bars guarantee a minimum of 60 hot plugs of a 500W chassis powered by 12Vdc in the same slot, before the bus bar plating deteriorates (for example, 6 chassis swaps are done in one year on same slot, over 10 years, with a new chassis every year). The nickel plating on the bus bars is minimally 10u thick and is normally made with an electro-plating process.

Hardened plating is recommended to extend the life of the bus bars; alternative plating processes are under investigation.

5.2 Power Shelf and Battery Backup Options

Online power and backup power can be provided by at least two power shelf configurations, called v1 and v2 here.

5.2.1 V1 Power Shelf

This 6+1 redundant shelf is nominally 3xOpenU high and requires the Open Compute Project battery cabinet to engage backup functions. The power modules deployed in this shelf are a variation of the Open Compute Project 700W-SH power supply, repackaged as a hot swappable module with handle, and adding more features like a power module failure signal and output OR-ing devices. The opto-isolated failure signals are daisy-chained between all three power shelves in each column to allow for reporting of a power module failure, identifying the shelf where the failure occurred. The failure is reported to the rack switch through a custom digital box sitting on top of the rack, with one Ethernet RJ45 output. The box has several inputs and types, serving multiple possible configurations and compatible for future uses.

The PDUs (one per column) distribute both AC and DC power.

One Open Compute Project battery cabinet sits in between a pair of triplets in the data center aisle, providing 48Vdc backup voltage used in the event of an AC power outage. The cabinet is normally offline and at system level exceeds 99.75% of equivalent AC UPS efficiency. For more information, see the *Open Compute Project Battery Cabinet Hardware v1.0* specification.

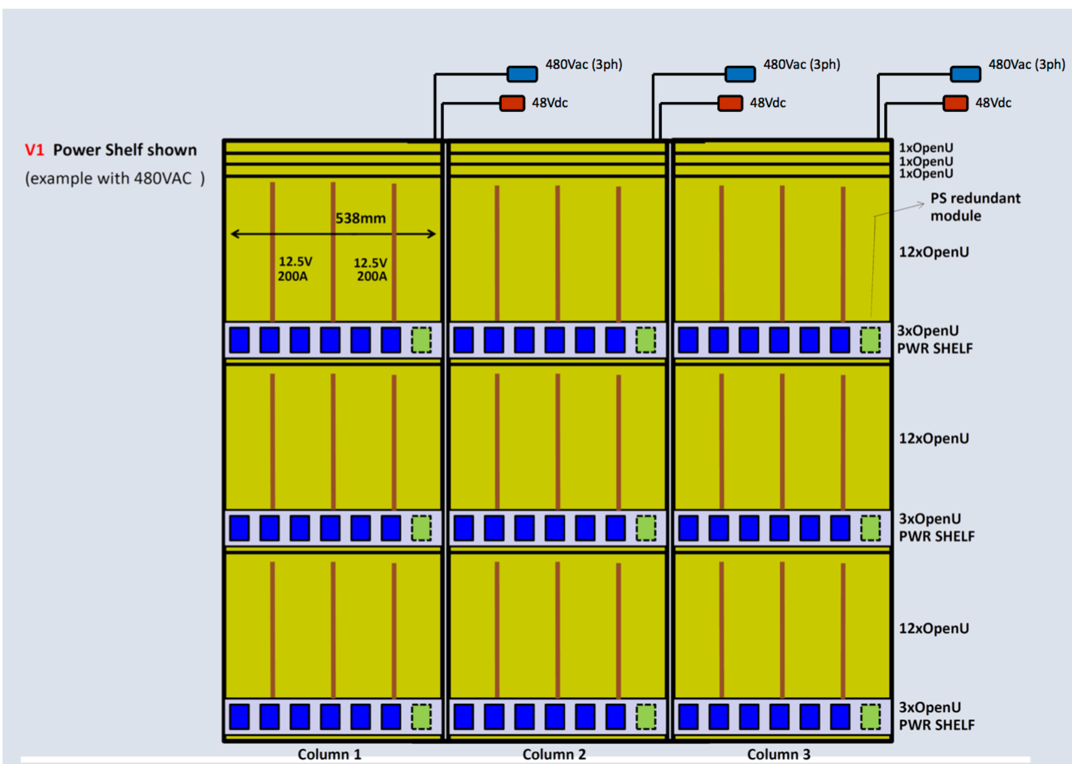

Figure 4 Triplet with Nine 3xOpenU V1 Power Shelves (OCP Battery Cabinet Not Shown)

5.2.2 V2 Power Shelf

This N+1 redundant shelf is nominally 2xOpenU (or 1xOpenU) high with a 1xOpenU battery backup unit (BBU) module normally installed underneath. The Open Compute Project battery cabinet is no longer needed as the BBU normally includes its own high-density lithium-ion battery pack. Each BBU module may also be placed above the 2U/1U power shelf to directly power the bus bar behind it, and may adopt any backup topologies. The BBU is N+1 redundant, the battery pack is also N+1 redundant. A battery monitoring and testing system can be included, and data is sent out via a digital bus. In the v2 scheme, the PDUs distribute only AC because BBUs with batteries are included in the power zones.

To account for power module failure signals and sophisticated power monitoring, a digital bus with more complete reporting on the power shelf functionality and health may be defined as needed for the V2 power solution, and it may or may not include the BBU reporting in the same bus.

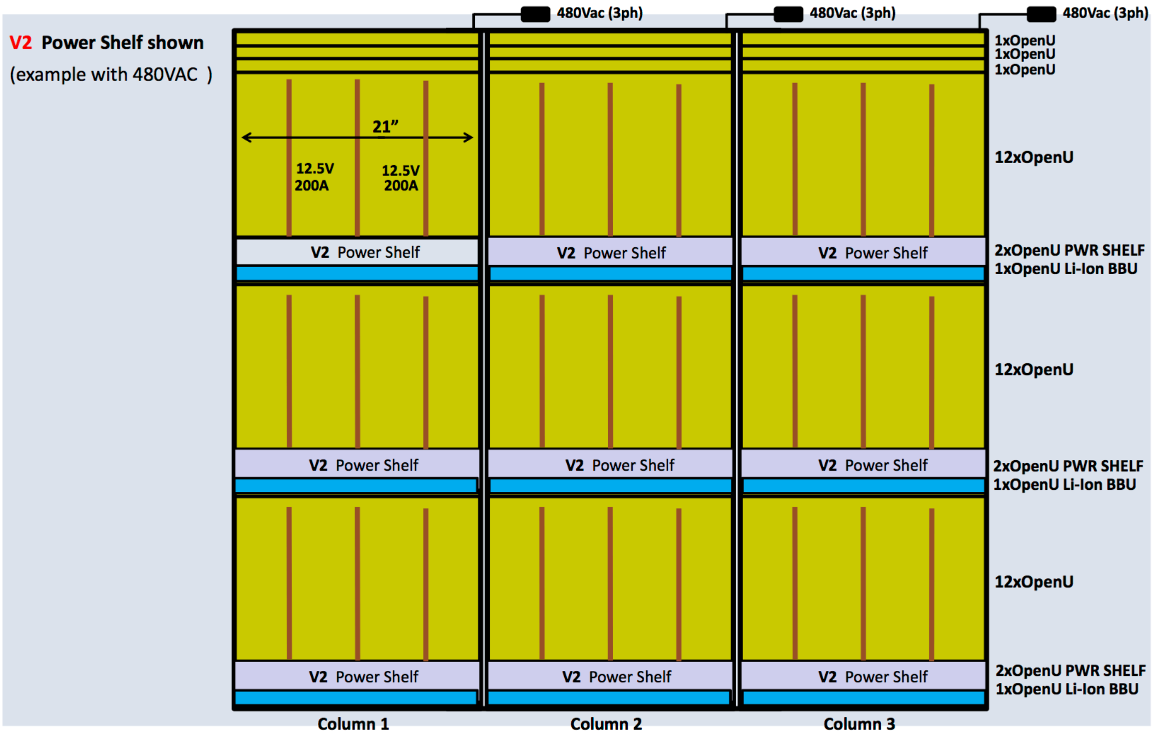

Figure 5 Open Rack Triplet with 2xOpenU Power Shelf and 1xOpenU BBU

5.3 HVDC Input Voltage Option

The Open Rack may be powered by High Voltage DC input (HVDC) instead.

An AC power shelf solution, after some customizations (such as input connector, circuitry adjustment, safety and regulatory compliance, and so forth), may support HVDC input directly. Otherwise a custom power shelf can be designed for HVDC input only. Typical voltage levels are in the range of 360Vdc ~ 400Vdc. This option gives the ability to power the rack in a data center environment with HVDC distribution or when the DC power is generated on site (for example, by fuel cell or other renewable energy source). An HVDC source is normally already backed up by batteries. Otherwise a BBU unit may still be included in the power zone.

6 Power Distribution Units (PDU)

An AC PDU may be required when there is more than one power zone in each column, as described in this example implementation. Alternately, the power shelves may be reached directly with a power cord, either a single-phase or a three-phase plug.

A DC PDU is required when the Open Compute battery backup cabinet is used as a backup power unit, as described in the v1 implementation, where the AC and DC are routed within the same PDU component, called a vertical PDU (vPDU). The PDU can be installed horizontally, known as an hPDU (see Figure 8). But an hPDU failure would cause the shutdown of 1/3 of each of the three logical columns (a logical column is a column in the triplet with its own dedicated IT switch), making this is a less desirable configuration than the vPDU, where the failure would cause the shutdown of just one logical column controlled by one IT switch.

The AC three-phase voltage is distributed Line-to-Neutral to each of the three single-phase power shelves in the column; for example, 480Vac nominal three-phase is distributed to three 277Vac single-phase feeds. The PDU also includes six AC sockets (one pair per phase protected by a 10A fuse per pair) that can be used for redundant power to the rack switches, or to power other devices. The DC voltage is distributed within the same PDU enclosure equally to the three power shelves in the column (48Vdc nominal). The DC voltage is used solely for backup power functions, and so it is normally offline; that is, the voltage is present but no current is delivered, until a power grid AC outage occurs.

6.1 Suggested PDU Dimensions and Configurations

The PDU for the v1 implementation is 50.8mm wide, to match the square tubing of the rack, and up to 101.6mm deep (AC and DC distribution in the same PDU) or up to 50.8mm deep (AC distribution only).

PDU height correlates to the rack height.

The AC voltage distribution in this example implementation is as follows:

- Power zone 1 is powered by 277V Line 1 to Neutral
- Power zone 2 is powered by 277V Line 2 to Neutral
- Power zone 3 is powered by 277V Line 3 to Neutral

The 277V value refers to the usage in the US, with three-phase power at 480Vac. For three-phase power at 400Vac, the single-phase voltage would be 230V.

For a power shelf solution using true three-phase power with or without Neutral, the AC PDU would distribute the same identical three-phase voltage to every power zone (4 or 5 wires, including GND conductor). Therefore a three-phase solution at the power shelf level is another option. This approach makes more sense when a high power density solution is needed, and in this way a three-phase balance is easily achieved without regard to the configuration of the load in the power zones. It also avoids usage of the Neutral conductor. However, a three-phase solution at the power shelf level that includes the Neutral conductor may still be considered, since it has other advantages.

In general, all the PDUs are installed at the rear of the rack and are flush with the rear of the rack.

AC and/or DC outputs from the PDUs are normally dongle cables terminated with connectors that must use strain relief or molding, and with the corresponding counterparts installed in the rear panel of the power shelf.

7 Implementation Details

This section describes some details of the v1 Open Rack implementation. It is a triplet rack using three V1 power shelves in each column, for three power zones per column. Three PDUs distributing AC and DC voltage in the same enclosure are used because the power shelf is single phase and includes a backup converter. The Open Compute Project battery cabinet is used as the energy backup unit.

Figure 6 Triplet Rack with PDUs (PDUs Are Flush to the Rear of the Rack in the Actual Implementation)

The power shelves are powered with single-phase (nominal three-phase voltage range 340Vac-480Vac RMS) using a single-phase Line-to-Neutral distribution with a PDU. In the US, the Open Rack uses 480Vac three-phase with a single-phase 277Vac distribution to each power zone. The power shelf is a single-phase power converter rated 200–277Vac nominal, supports a voltage range from 180Vac to 290Vac for worldwide compatibility, and supports backup power using 48Vdc nominal input (38Vdc to 58Vdc).

Each power shelf supports up to 4.2KW at 12.5Vdc output, for a maximum of 13.3KVA of AC power (277Vac x 16A x 3 phases) in a single column when using a US standard NEMA

L22-20P plug. The AC plugs are normally derated 20%; the L22-20P becomes a 16A plug. These ratings factor in an efficiency of ~ 0.94 for an OCP power supply module (for example, 4.2KW = 277Vac x 16A x 0.94 efficiency) and nearly unity power factor (PF) at full load. Higher power per column may be achieved by hard wiring the AC power cord to the grid or by using a higher power plugs (like NEMA L22-30P, or industrial plugs like the pin and sleeve type).

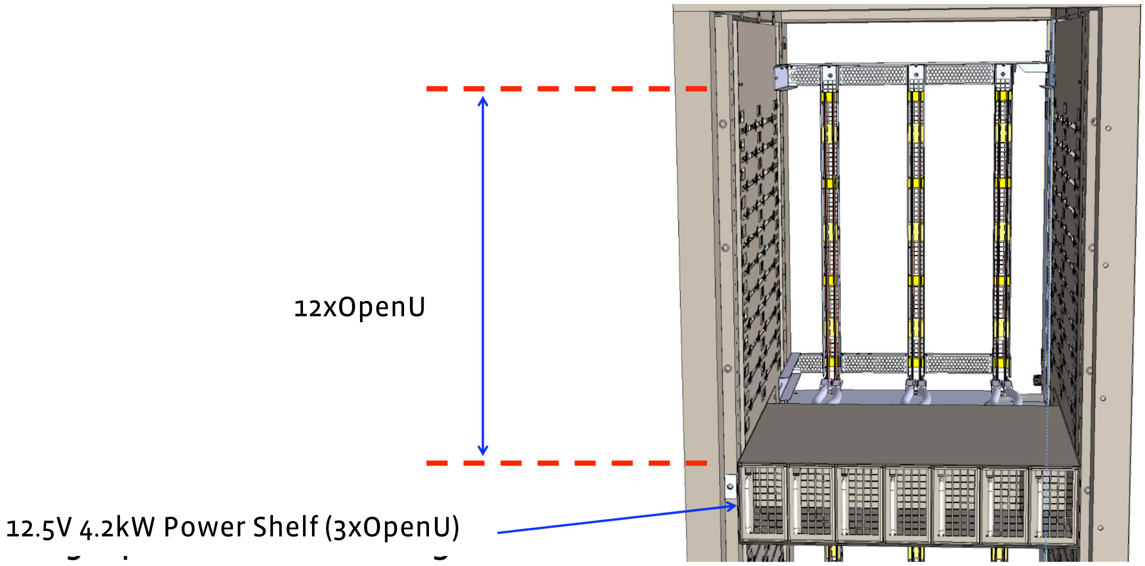

Figure 7 Power Zone: 3xOpenU Power Shelf, 12.5V, 4.2kW Power

The example rack uses three-phase voltage distributed into single phases for flexibility and cost effectiveness. Each column in a triplet gets one phase for each power zone. To achieve three-phase load balance, each of the three power zone loads are configured identically to the other two.

The rack supports backup power functionality, with no interruption of service in case of an AC outage. Since the OCP battery cabinet is used, the 48Vdc backup voltage is connected to the V1 power shelf by means of high-current blade connectors APP SB175, where one connector supports three power zones within the same PDU.

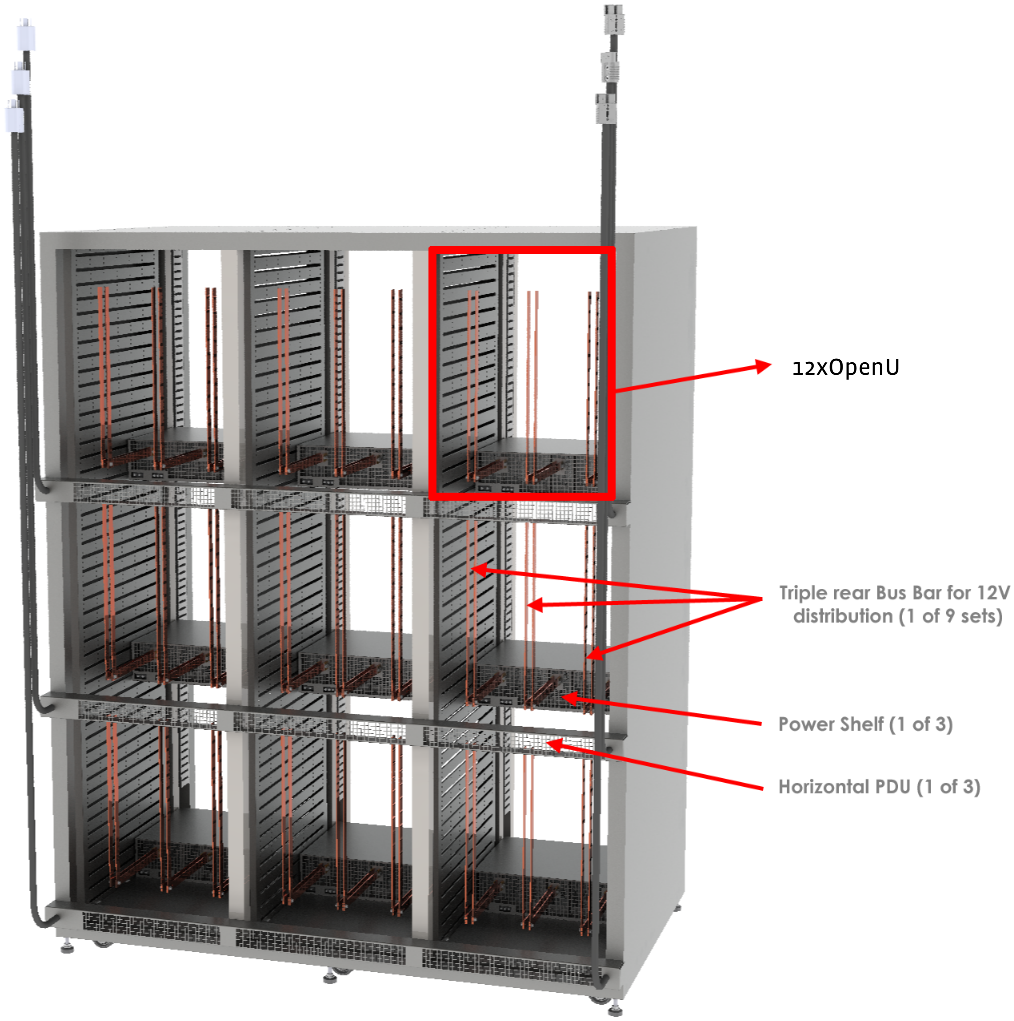

Figure 8 Open Rack Triplet Rear View (Horizontal PDUs Shown)

8 Conclusion

This specification is a work in progress. To offer feedback or contribute, join the mailing list at http://lists.opencompute.com/mailman/listinfo/opencompute-openrack.