Reimagining Rack PDUs Through Modular Design

Introduction

Since the mid-2000s, the data center industry has been taken with the idea of modularity. The adjective "modular" has been applied to various products that end users might consume—from containerized data centers to electrical and mechanical products that provide essential power and cooling to IT hardware. A quick internet search of "modular data center" finds millions of webpages competing for this search term.

Modularity was (and is still) presented with the joint value propositions of economy and efficiency. The tagline is often: buy what you need when you need it. The virtue of modularity is also seen in service and maintenance, as failed components can often be replaced without a costly wholesale exchange of an entire device.

Modularity in Power Backup and Power Distribution

During the same decade, we saw the introduction and adoption of products that users could bolt onto cabinet rows and bolster capacity as they added IT resources. As such, models of air conditioners and uninterruptible power supplies (UPS) were no longer limited to big-box form factors. They had an aesthetic quality and footprint that closely matched that of an equipment rack. At the same time, the larger facility UPS began to incorporate modular ideas that allowed for easier upgrades, expansions and redundancy planning.

Modular UPS first came into fashion targeting small-to-medium applications that would typically top out around 100kVA. Vendors would provide a fixed amount of capacity to start (e.g., 20kVA), but included built-in headroom with the UPS, batteries and distribution to grow to 100kVA as the business filled the data center with more equipment.

Another example of power modularity is the overhead busway with tap-off boxes. These products became prominent in the data center in the mid-to-late 2000s, bringing a track lighting-type simplicity to rack power feeds. As long as the busway was sized appropriately, a user could support a bevy of power configurations in a given row of equipment cabinets; a single-phase 20-amp feed could be deployed next to a three-phase 60 amp feed. If server churn was high, there was no need to worry about the laborious work of adding or changing circuit breakers in a panelboard; the old and new tap-off boxes could be replaced by data center personnel in a matter of minutes. Figure A below demonstrates an example of power modularity through tap-off boxes on an overhead busway.

The rack power distribution unit (rPDU) is the last mile in the data center power chain and the product itself is no stranger to modular variants. Certain vendors have turned the rPDU chassis into a touch-safe busbar where the customer can combine different modules of breakers, receptacles and intelligence (boards, network ports) that lock into place. This erector-set approach to the rPDU results in a highly bespoke product with a variety of outlet types to suit any IT deployment and one that can be quickly changed in the event of server refresh.
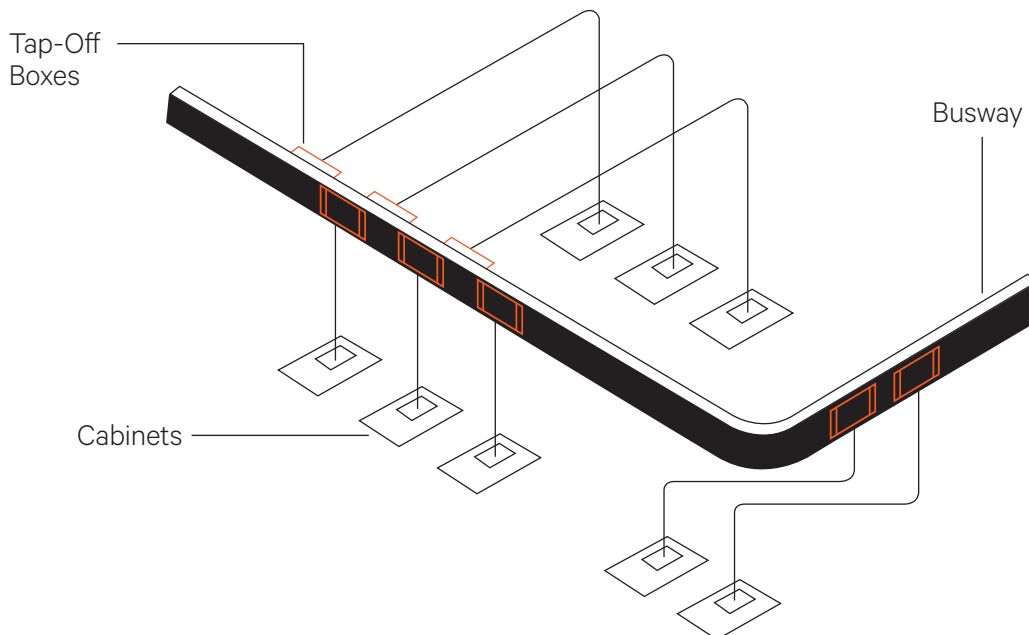

Figure A: Overhead Busway with Tap-Off Boxes Example

Modularity & Component Anatomy Intelligence

Modularity is about component parts—understanding what has normally been hard-coded and then engineering those pieces to be flexible and interchangeable instead. This philosophy is expressed by the **Open Compute Project (OCP)** consortium who, in their Open Rack charter, write:

"The intent is to enable what we term component disaggregation [where] compute components can be swapped and expire according to their own life cycles without replacing the entire server for an upgrade or repair."

A rack PDU, made in this mold and embodying this disaggregation, would step outside the conventional market model for rack PDUs.

Rack PDUs have historically been segmented by technology into the lexicon of basic, locally metered, input metered, outlet switched, outlet monitored or outlet switched and monitored. As users progress from one segment to the next, functionality increases and so does cost. But, to this point, the segments

have yet to be bridged by technology; the upgrade path has been rip and replace—much like the hardware model that the OCP is trying to counter. An alternative to this course is a rPDU with hot swappable intelligence—modules that enable feature upgrades while retaining the same rPDU chassis.

This path realizes that basic rPDU users today may have a different outlook in two years; the business need may be more complex and IT may require more visibility. Similarly, consumers with switched rPDUs may want to implement an efficiency program that requires more granular measurements at the outlets. Through hot-swappable intelligence, both users achieve a degree of future-proofing; see example in Figure B below. The rPDU becomes an investment that can be repurposed as power management priorities change.

And, as rare as rPDU failures are, the root cause of such events often traces back to these embedded systems. In these infrequent cases, the likely culprit can be replaced without any rPDU downtime and visibility to rPDU information can be quickly restored.
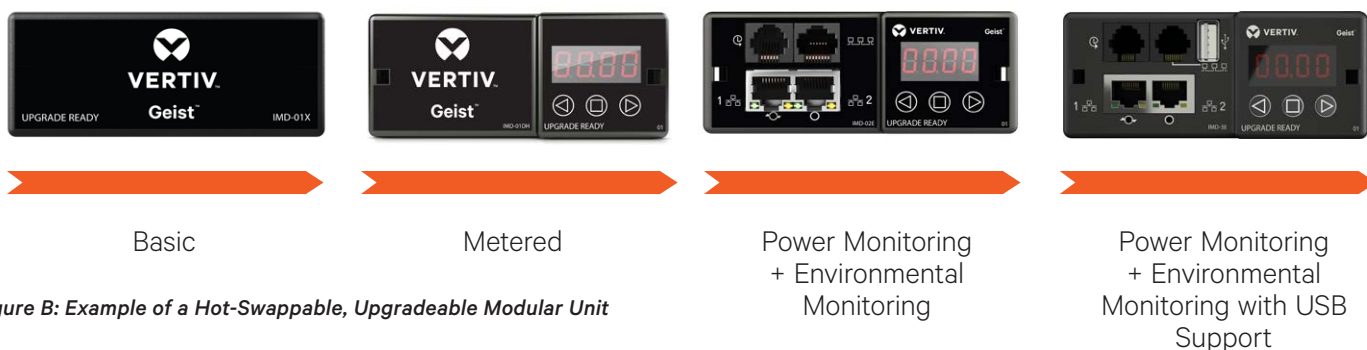

| Basic | Metered | Power Monitoring + Environmental Monitoring | Power Monitoring + Environmental Monitoring with USB Support |

Figure B: Example of a Hot-Swappable, Upgradeable Modular Unit

Chassis & Cord

Beyond intelligence (sensors and chips), a rack PDU is composed of various component parts—a metal chassis, conductors, receptacles, overcurrent protectors, a cord and a plug. As mentioned, these components have typically been hardwired as a complete product; however, interesting possibilities are created if some of the components are disaggregated from the chassis.

Most data center operators deploy a variety of rack PDUs based on hardware applications or deployment locations—whether the hardware resides, for example, in an on premise (company-owned) facility or a leased location. Some large operators in North America have realized the benefits of operating at the typical international 400V standard in their company-owned facilities. These same operators may also

maintain points-of-presence or other production gear at North American colocation facilities. These off-premise sites are more likely to run the prevalent 120V/208V power architecture. In these cases, managers must make decisions about inventories, sourcing and supply chain optimization across multiple SKUs—decisions that may be further complicated by an international portfolio and diverse power standards per country of operation.

- **Typical North America Voltage Possibilities:**
  120, 208, 208 3-phase, 400 3-phase

- **Typical North America Amperage Possibilities:**
  15, 20, 30, 35, 40, 50, 60, 80, 120

- **Typical International Voltage Possibilities:**
  100, 200, 230, 400 3-phase

- **Typical International Amperage Possibilities:**
  10, 13, 16, 32, 40, 48, 63

To counter this complexity, some organizations have found value in standardizing on a rPDU chassis that is complete with receptacles, conductors and breakers but uses specialized connectors with a cord that is removable and interchangeable. In this scheme, the physical rPDU is "universal" across all voltage and amperage considerations and is defined entirely by this modular cord and plug—of which there could be several possibilities. Thus, the same physical rPDU could be deployed across the world, while individual sites would only manage their cable inventories.

From the facilities side of the data center, the universal rPDU can offer a simpler provisioning of rack power. As figure D depicts, the typical connectors are no longer external to the rPDU; they are now a part of the physical chassis. The "facility end" becomes one connection to an overhead busway, removing the intermediary devices.

This universal concept is the converse of how most rPDUs are made and sold today. As such, it requires some unique components. The most unique being the cable which contains a standard plug (e.g., a 50 amp CS8365), a specialized male

connector and a matching female connector permanently affixed to the rPDU. The conductors are uniquely spliced within the cable so that, once connected to the female end, the downstream rPDU components (breakers and receptacles) become wired either line-to-line (208V 3 phase) or line-to-neutral (400V 3 phase).

The physical rPDU would see little difference from a standard model, aside from oversized wiring between the female connector and the breakers' input terminals to accommodate the highest possible amperage.
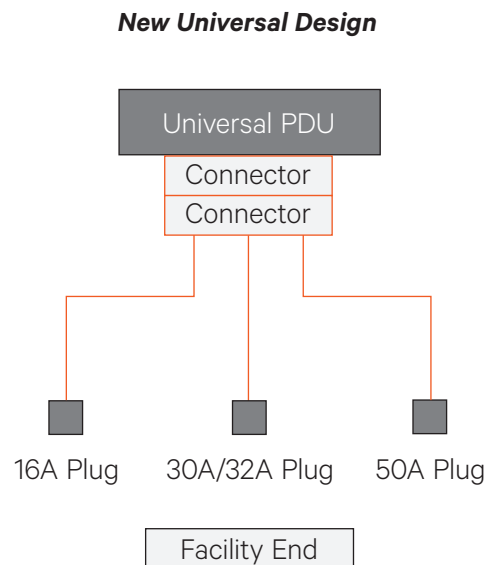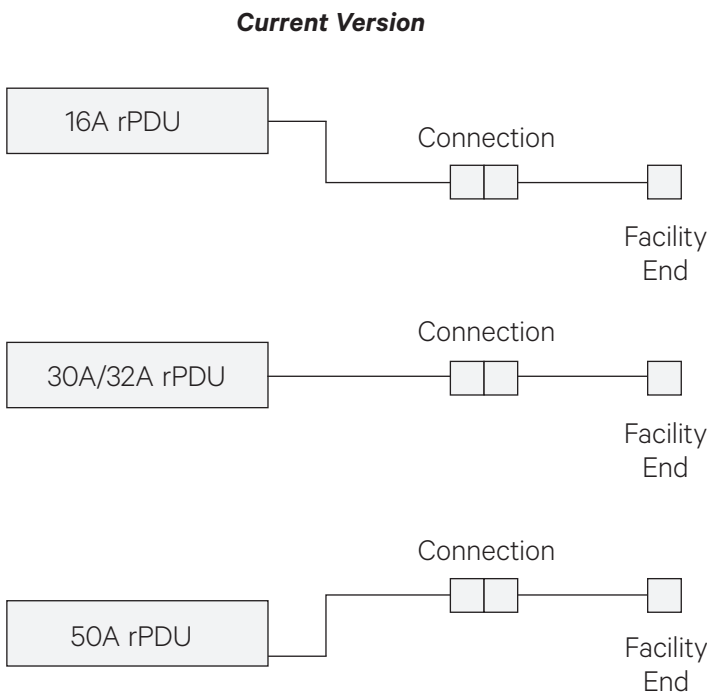
Figure C: Standard rPDU Versions vs. Universal PDU Version
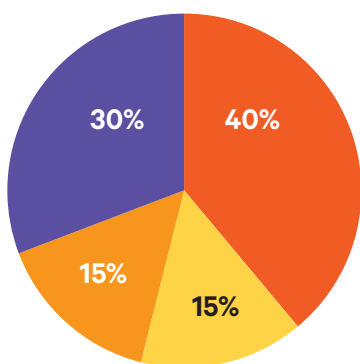
"Bare-Metal" rPDU

As shrinking hardware footprints have allowed users to pack more IT in a single server cabinet, these users have demanded the same from their rPDU suppliers: maximizing receptacles in the smallest possible form factor. Input amperage, input voltage, receptacle counts and level of intelligence all influence a rPDU's size. These factors can be a challenge for a manufacturer to build a high-amperage, outlet-dense, fully instrumented rPDU when confined to a typical 72" x 2" x 2" form factor. If such a design is realized, it may come at the expense of the best outlet spacing, as the rPDU supplier must balance usability, component parts and chassis real estate all in tandem.

All told, this example is a perfect one to revisit modularity and disaggregation, as the rack PDU is often seen as the demarcation point between Facilities and IT teams within a data center.

Figure E details component costs for an intelligent rPDU. There are four primary categories, the costliest being physical components like receptacles, wire and chassis. In searching for efficiencies, we must consider the interplay between these parts:

- What makes the most sense - to keep the components together or break them apart?

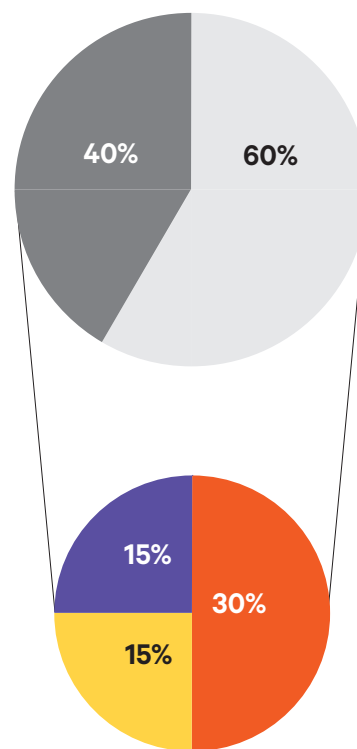- Are there particular components that are more allied to IT or Facilities?

- What value can a disaggregated design bring to each business unit?

Asking these questions, we can re-draw the pie chart (Figure F) into two main categories that align component parts to each business unit.

Then we can ask a simple question: what if the most expensive part of the product became a permanent fixture in the facility installation? And, to tie back to OCP ideals, what if these components can be swapped and expire according to their own life cycles?

The lifecycle for a rack PDU can be tied to corporate technology refreshes, which happen on the average of three to five years. Following OCP's ideal, the components' useful lives are much longer. If a data center can achieve long-lasting utility from 60% of its initial rPDU investment, there are compelling financial benefits to the organization.

Intelligent Rack PDU Cost Breakdown - Reformatted

New Universal Design

- 🟧 Receptacles, Wire, and Chassis
- 🟨 Plug and Cables
- 🟧 Overcurrent Protection
- 🟪 Monitoring

Figure D: rPDU Cost Breakdown

- ⬛ Receptacles, Wire, and Chassis
- 🟨 Plug and Cables
- 🟪 Overcurrent Protection
- 🟧 Monitoring

Figure E: rPDU Cost Breakdown - Reformatted

This concept presents another spin on a universal PDU. Instead of a busway with discrete tapoff boxes (e.g., an L21-30 rack next to an L6-30 rack) users deploy the same tapoff box (a "facility" rack PDU) with the same connecting cable to each cabinet. This facility box is preconfigured with the following components:

Overcurrnet protection containing six branch circuits for high density possibilities

—

Monitoring hardware

—

Ethernet ports for network interfaces that intergrate with IT and Facility tools

The facility output would be hardwired to a specialized cable and male connector that would facilitate any number of rack kW possibilities within its connected cabinet.

Figure F: Facilities rPDU

The disaggregation of components results in an "IT" rPDU with optimized form factor: low profile and outlet dense, where receptacle positions are closely matched to rack units and phase balancing is easy to accomplish. During the next technology refresh where new, fully populated racks are rolled-in by systems integrators, the legacy racks, hardware and IT rPDUs can be decommissioned. The new IT rPDUs, purchased at less capital expense, can easily connect to the existing facility rPDUs via the remaining cords.

This modular approach also offers additional functionality, centered on risk prevention and monitoring.

Overcurrent protectors in rack PDUs have not traditionally been subject to the selective coordination rules that apply to the rest of the data center's power chain (i.e., rack PDUs, like the IT gear they power, have historically been classified as "utilization equipment"). With breakers now housed in the Facility rPDU and tied to the busway, they will be considered part of the primary distribution path and under the purview of selective coordination—a practice that should bring solace to facility managers and electrical designers.

The Facility rPDU includes further modularity within its intelligence card. The same technology mentioned earlier providing the consumer with a built-in upgrade path that is non-invasive and hot-swappable.

Figure H: Complete Installation

Figure G: IT rPDU

Conclusion

The data center's modular evolution has touched nearly all products in the white space and many outside of it. While rack PDUs are not exempt from this trend, there has not been considerable historical demand for the product to take a more modular turn. But market conditions are changing. Large, influential data center operators are breaking down the facility and its components into bare metal, open source components, valuing utility and flexibility over other considerations. The rack PDU models herein embody this theme, presenting products that can be easily upgraded, continually repurposed and disaggregated so that each component part can maximize its value.

About the Author

Justin Blumling
Offiering Manager, Vertiv

Justin Blumling is an offering manager for Vertiv's global rack PDU (rPDU) line of business. He has spent his career in the data center industry working for both OEMs and VARs in sales engineering, marketing, and project management roles. He holds a Masters of Business Administration degree and is a US Department of Energy-certified Data Center Energy Practitioner.

Vertiv.com | **Vertiv Headquarters,** 1050 Dearborn Drive, Columbus, OH, 43085, USA

SL-70502 (R02/19)