DATA CENTER TRANSFORMATION: THE IMPACT OF EMERGING POWER ARCHITECTURES ON TODAY'S DATA CENTERS

Entering the Cloud Generation

The unprecedented capacity requirements of large social media, search, colocation and cloud companies is driving massive investments in data center development. The organizations creating this capacity are continually experimenting with new technologies and designs to push the limits of data center performance while driving down costs.

Their scale, individually and collectively, has the potential to drive significant change in the data center industry. It allows them to work with their vendors on custom designs and solutions specific to their needs, bringing new solutions to the market. It also gives them the ability to test multiple designs simultaneously, often within the same facility, to determine which best meets their demands for flexibility, speed and reliability at the lowest cost. While some of the designs and technologies that emerge from this development will be specific to the largest data centers, others will have broad applicability.

In fact, these large developers and operators are creating a new generation of data center that continue to change the landscape of the industry. Just as the mainframe generation gave way to the client/server generation, the client/server generation is now being replaced by the cloud generation (Figure 1). This generation of data centers is marked by a philosophy of deploying only "what is needed, when it's needed" to support a particular application or set of applications fast and efficiently. To support their goal they are evaluating new power system architectures and the best location for backup power within the data center (Figure 2.)

In this report, we'll overview several alternative power configurations that improve overall cost and deployment speed while providing the availability levels required for this new generation.

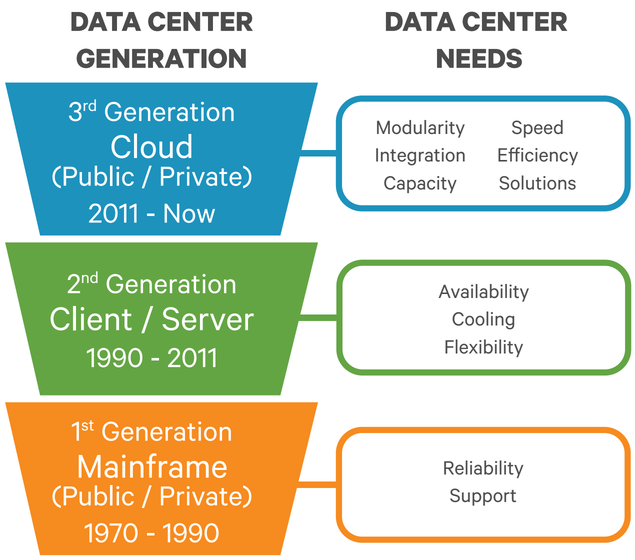

Figure 1: *Figure 1. Requirements have changed as the data center has evolved from the client / server to the cloud generation.*

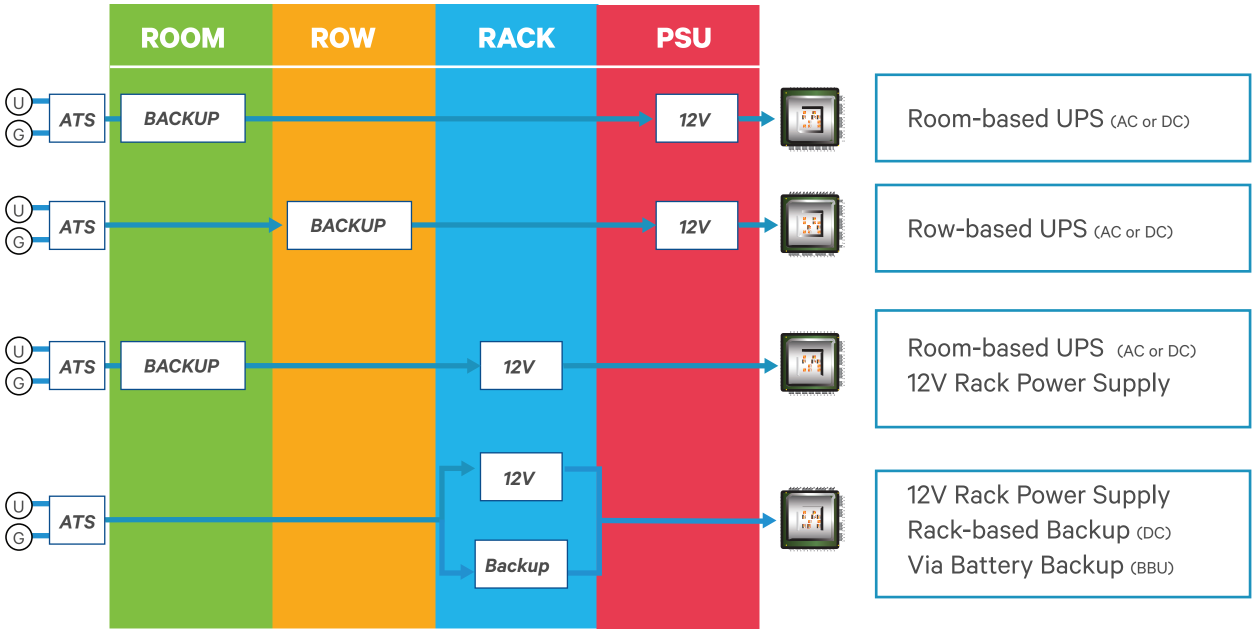

Figure 2: *New generation data centers are evaluating multiple options for deploying power backup.*

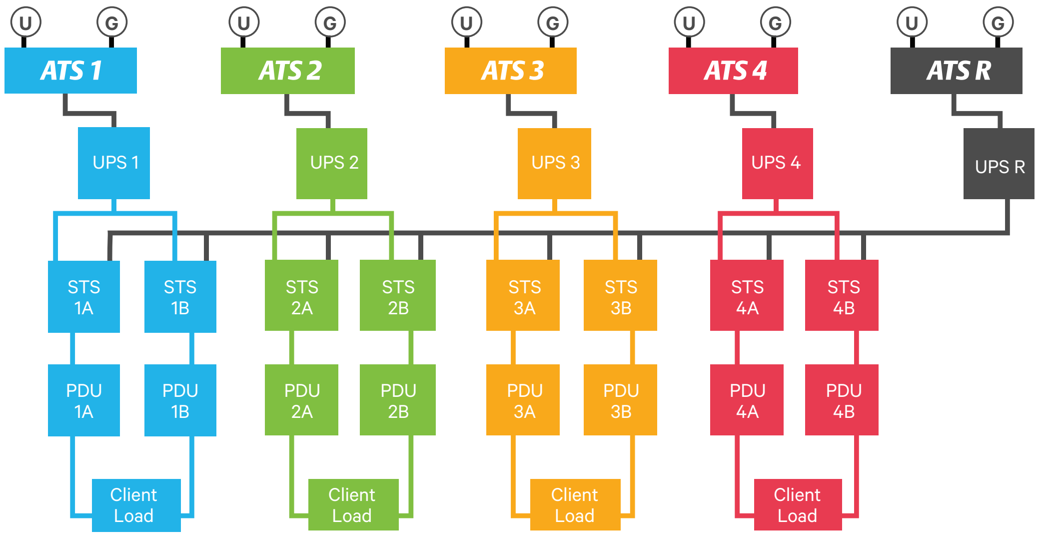

Figure 3: *A reserve power system which uses four UPS modules to support the load with one in reserve.*

Using the Reserve Bus Architecture to Streamline Redundancy

The 2N or 2N+1 dual-bus architecture has historically been the choice of high-availability data centers. When properly designed, these architectures eliminate single points of failure in the critical power system and allow maintenance to be performed on any component while continuing to power the load.

However, in today's environment where the need to optimize capital efficiency and resource utilization is paramount, this level of redundancy is becoming more difficult to justify. Increasingly, the 2N + 1 or 2N dual bus architecture is being replaced by various reserve architectures pioneered in large colocation facilities.

The basic reserve architecture creates a redundant architecture, while maintaining fault tolerance and concurrent maintainability through the use of static transfer switches (STS). The STS allows a redundant UPS system to be brought online to pick up the load from any one of multiple UPS systems in the event of failure or maintenance. Downstream from the STS units, the power distribution system can be similar in design to that of a 2N dual-bus architecture.

This deployment does complicate maintenance and load deployment compared to a traditional 2N architecture, but the economic benefits are compelling. Consider a 2N + 1 architecture consisting of six 1100 kW UPS modules. If the modules are sized to 110 percent of maximum load, the system is capable of supporting 2000 kW. Shifting to a shared reserve architecture, in which five of the modules are supporting the load with one reserve module, the same UPS capacity can support 5000 kW. High reliability reserve architectures, such as break one - fix one, can also be achieved.

Variations of the reserve configuration can be considered. The primary difference in the configurations rests with how the client loads achieve power redundancy: either sharing a reserve system as shown in Figure 3, dedicating the reserve system to high-priority clients or accessing unused capacity across multiple UPS modules to create the reserve.

Using the Reserve Architecture to Mimic 2N

In the dedicated reserve architecture (Figure 4) higher levels of availability can be supported by directly tying the reserve power to a specific client or application. This dedicated reserve ensures redundant capacity is allocated to support specific UPS loads. Colocation providers may use dedicated reserve modules to provide 2N backup capacity for customers requiring higher SLAs.

Alternately, two reserve modules can be shared across multiple primary modules in a configuration that is commonly referred as “eight to make six” or “ten to make eight.” With this configuration, any module can be taken offline for service while maintaining redundancy across the system.

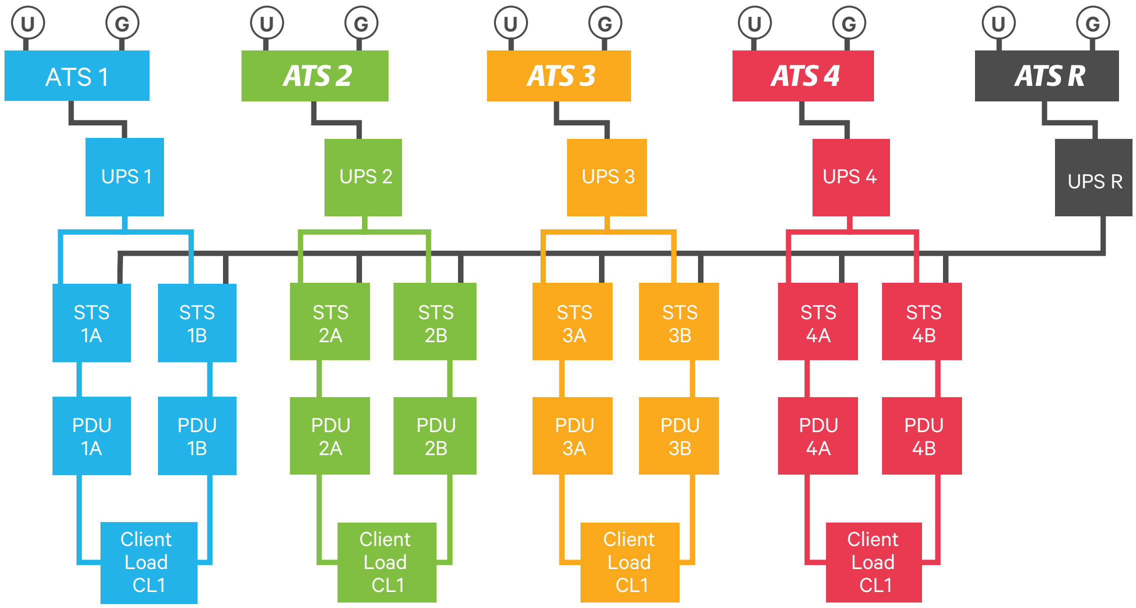

Figure 4: *A dedicated reserve power system is provided to client load CL1.*

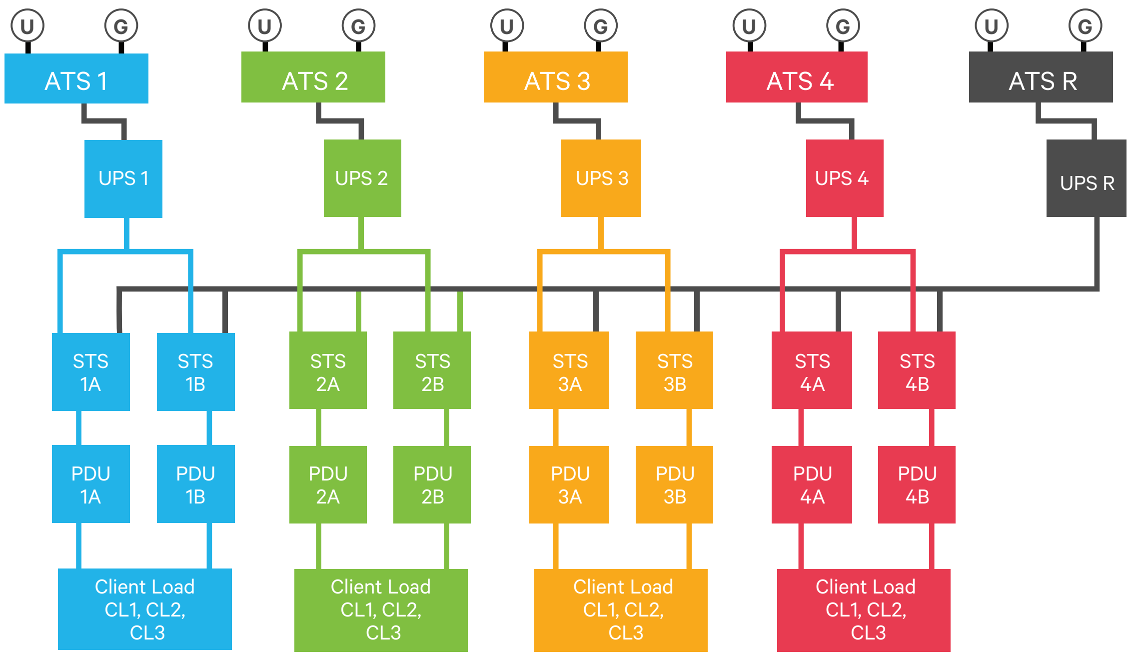

Figure 5: *Shared Reserve differs from the dedicated reserve by the applied support plan of client loads.*

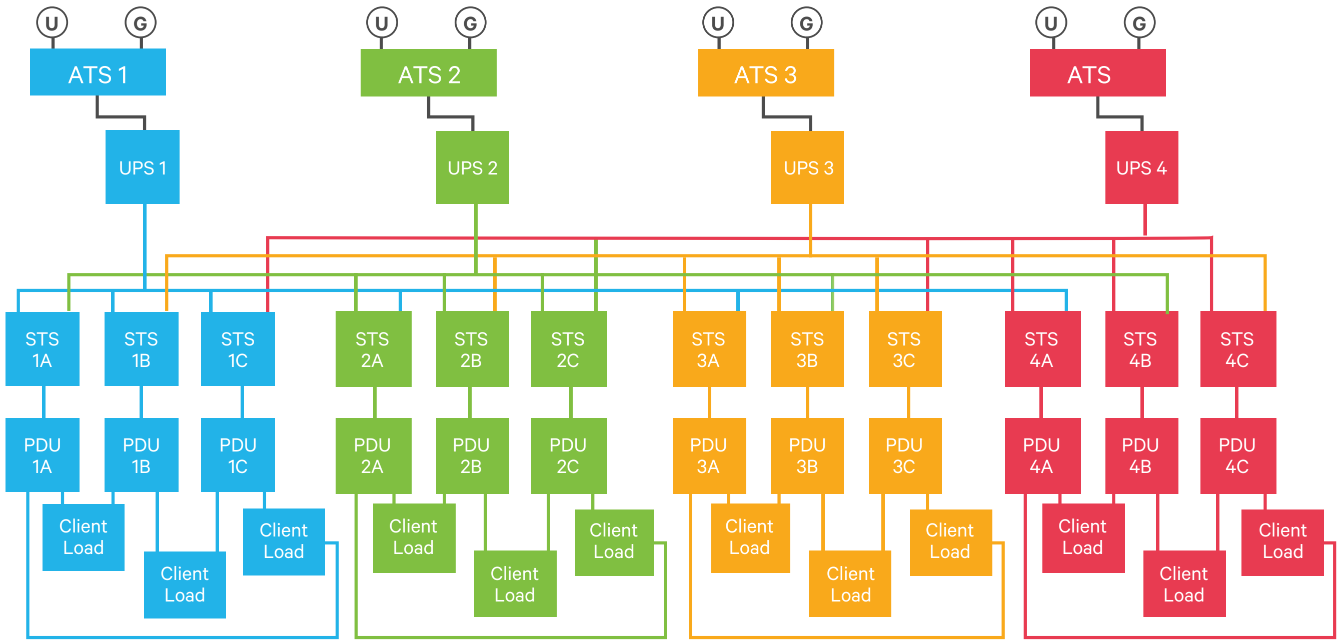

Figure 6: *A Distributed Reserve utilizes available within the existing UPS system.*

In a shared reserve configuration (Figure 5), the reserve power system is shared across more than one customer or application.

Utilization is slightly lower in a dedicated reserve than can be achieved with a single reserve module but still higher than is possible in the traditional 2N architecture. Plus, new capacity can be supported through the addition of one module rather than two as would be required with a 2N system.

A distributed reserve system can be seen in Figure 6. In this case, the reserve power system is achieved by utilizing the unused capacity within the UPS system modules. Here the distributed reserve is allocated across the UPS loads either on a first-come first-served or policy basis.

The reserve power configuration, whether shared, dedicated or distributed, offers significant flexibility in the quest for efficiency, speed and availability and these configurations have applicability within both collocation and enterprise applications.

The Importance of a Critical Power Management System

A critical power management system (CPMS) is highly recommended for any reserve system implementation. It proactively manages loads and capacities to maximize reserve system utilization, while performing successful transfer procedures that prevent overloading any reserve system module. The CPMS provides optimal power management across the power chain and unifies control and reporting.

Enhancing Flexibility with Rack-Based Power Protection

For the developers of many large data centers, speed of deployment has risen to the top of the list of design criteria. They need to bring on capacity quickly and incrementally without compromising capital efficiency. One way to accomplish that is by driving power protection to the row and ultimately to the rack (Figure 7), making the rack an autonomous unit that can be brought on line without adding to the load of a room- or aisle-based power protection system.

The simple approach to implement this scheme would be to place UPS systems in each rack, but that doesn't fit with the design philosophy of deploying only what is needed in a modular, integrated form. Developers can now deploy rack-based power systems to energize DC-powered servers inspired by the Open Compute Project. This centralized rack-based power system comprises rectifiers for main power (replacing the AC/DC power supply traditionally embedded in an AC-powered server) supported by lithium

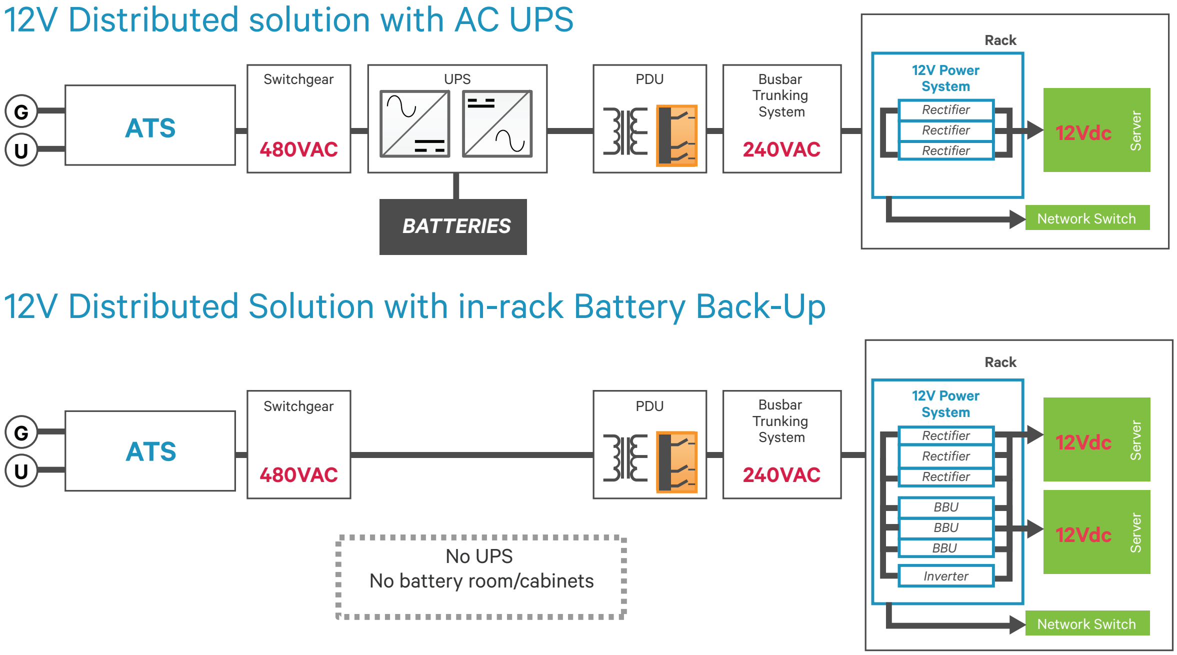

Figure 7: *Displays the power path of a rack-based power supply with row UPS and a rack-based power backup.*

ion batteries for power protection (substituting for the UPS). The rectifiers receive 480V or 240V unconditioned AC power and convert it to 12V DC power for use by the servers. In the event of a power interruption, the lithium ion batteries provides short-term ride through of 12V DC power. The result is a relatively efficient and economical backup power strategy that provides the ultimate in flexibility by enabling capacity to be added one rack at a time.

The maturation of lithium ion technology is a key enabler of this strategy as it provides a compact backup power source able to support short discharge times with a high discharge cycle count, high power density and the ability to operate in the increasingly high temperatures that exist in this new generation of data center.

Embracing Simplicity and Efficiency with High-Voltage DC Power

Convergence of voice and data has dictated that telecommunications providers become major data center developers. They bring a long history with 48V DC power, with its proven reliability and efficiency, to the traditional data center. However, 48V DC has not proven practical in the data center environment due to the challenges of distributing low voltage DC power.

High voltage DC power (Figure 8) brings the benefits of DC power to the data center while eliminating the high infrastructure costs associated with distributing lower voltages. Deployed at either the room or row level, the DC UPS converts AC utility power to 400V DC power through a bank of rectifiers in the UPS. The DC UPS is sized to withstand a failure of any rectifier without impacting operation, creating internal redundancy that eliminates the need for redundant configurations common in AC UPS systems.

DC power is then transmitted to rack power supply units (as in the rack-based power distribution strategy previously described) or sever power supplies which step down the DC power to voltages that can be used by components, or fed directly to the server motherboard thus eliminating the server or rack power supply. The promise of high-voltage DC is that it can simplify power system design and management, enhance scalability, and increase efficiency. DC distribution can be easily configured for any desired redundancy, including N+1, N+N, 2N or DC/AC hybrid configurations.

The biggest challenge facing DC power has been the immature supplier ecosystem, but if just a few of these new developers embrace DC power as they appear to be doing, their scale will create the demand that forces the ecosystem to be mature quickly.

380Vdc System

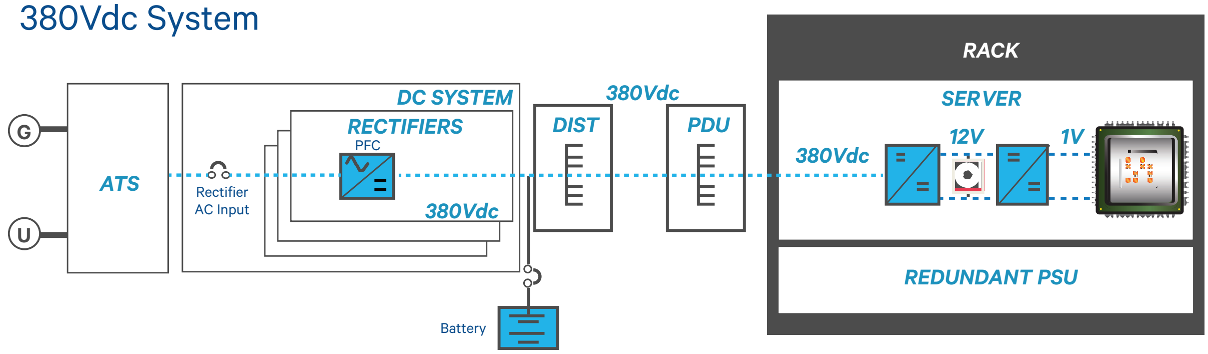

Figure 8: *Power is delivered to rack via high voltage DC*

The Trickle Down Theory of Data Center Innovation

More than a thousand megawatts of data center capacity will be developed in the next several years by a relatively few companies. Those companies are under intense pressure to build strong, ultra efficient data centers faster and cheaper than ever and they assessing every technology and practice with a critical eye in their efforts to accomplish that. In the critical power system, they are evaluating the best location for power backup—room, row or rack—are seeking to minimize hardware redundancy and are driving greater simplicity and integration in the power path.

Designers and users of these electrical architectures will want to effectively manage the total power stream within these configurations. That will necessitate the use of advanced critical power management systems to provide real-time monitoring and control of load capacity, switching, power quality and more.

This wave of development—and the innovations that emerge from it-- will bring new choices to organizational of all sizes not only in how they acquire capacity, but in how they deploy it and support it within their own facilities.

For more information on power configurations and technologies, please visit VertivCo.com/knowUPS

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

SL-24686 (R02-16)