

*Power Management Strategies for  
High-Density IT Facilities and Systems*

## **Executive Summary**

Increases in data center density and diversity are driving change in the power and cooling systems that business-critical servers and communications devices depend on for their performance and reliability.

Rising equipment densities often correlate with increased criticality as companies deploy new applications that increase business dependence on data center systems. At the same time, entire facilities, as well as individual racks, are supporting an escalating number of devices as server form factors continue to shrink.

This is creating the need for dynamic data center infrastructures. When critical infrastructure systems can respond to changes in density, capacity and availability created by new technology and changing business conditions, the result is greater operating flexibility, higher system availability and lower total costs.

In the area of critical power, a dynamic infrastructure must encompass the UPS system, the power distribution system and in-rack power management.

Redundancy within the UPS system can support higher availability and increased flexibility, but only if UPS modules are properly sized. Analysis of mean-time-between-failure (MTBF) rates indicates that single-bus UPS configurations requiring more than four UPS modules pose a reliability risk that will be unacceptable in many applications.

The growth in the number of devices that must be supported is driving change in the power distribution system. Power distribution is evolving from single-stage to two-stage designs to enable increased scalability, reduced cabling and more effective use of data center space.

Finally, individual racks are now supporting as much as 20 kW of technology, driving the need for in-rack power management. In-rack power distribution provides increased control and visibility of rack power while simplifying cable management and increasing rack airflow.

Integrating UPS, power distribution and in-rack power management strategies creates a power system infrastructure capable of achieving the availability and scalability today's data centers and network closets require.

## Introduction

Changes in the quantity, configuration and density of servers are reshaping the data center environment. Higher heat emissions are creating hot spots that threaten server reliability, while driving change in how data centers are cooled. This heat is created by increased power consumption, which impacts the UPS system and the power distribution system. Failing to design flexibility into these systems can significantly shorten data center life and threaten equipment availability.

Consider how power requirements for IT equipment racks have changed in the last 10 years. In 1996, a fully populated server rack could house 14 single-corded servers operating at 120 Volts. This rack consumed approximately 4 kW of power. By 2001, a fully populated rack could house 42 servers, which were likely to be dual-corded and operating at 208 Volts, single phase.

In a matter of five years, the number of power receptacles required to support a fully populated rack grew from 14 to 84, and total power consumption increased from 4 kW to almost 20 kW. Now, the emergence of blade servers is driving even more change. Today, a standard rack can house six dual-corded blade chassis operating at 208 Volts, single phase, with a power consumption of 24 kW.

This evolution has left data center managers dealing with rising power consumption, increased demand for circuits, and greater diversity across the facility. These challenges have created the need for a power infrastructure capable of adjusting to changes in the number of devices, the density of those devices and where those devices are located.

Such an infrastructure must encompass critical power management from the utility input to the point of use, including:

- Power availability as determined by the UPS type and configuration.
- Power distribution from the UPS to the rack.
- In-rack power management.

## Power Availability

The internal design of a UPS system determines the relationship between the UPS and incoming utility power and, ultimately, the effectiveness of the UPS at protecting against certain types of power disturbances. There are three types of UPS in use today: passive standby, line interactive and online double conversion. The online double conversion reverse transfer topology is the only one that protects against the full range of power disturbances and is recommended for applications that are currently mission-critical or expected to become mission-critical, which includes virtually all data centers. Selecting an online double-conversion UPS ensures that availability requirements will not outgrow UPS topology.

Using redundant UPS modules helps ensure appropriate levels of availability can be achieved. Depending on the configuration type, redundancy may also enable scalability. Table 1 shows a summary of the most common system configurations in use today.

Single-module systems are typically able to support availability levels of 99.99 percent, or less than one hour of unplanned downtime annually, and leave critical systems unprotected during UPS maintenance.

UPS redundancy adds availability and allows UPS modules to be serviced without impacting the quality of power to connected equipment. Single-bus systems with redundant UPS modules can support availability of 99.999 percent or higher, while dual-bus systems are designed to support continuous availability by eliminating single-points of failure between the UPS and the connected equipment.

The simplest approach to redundancy is a 1+1 system, in which each UPS module has the capacity to support all connected equipment. This configuration provides redundancy with the minimum number of UPS modules,

*Failing to design flexibility into these power systems can significantly shorten data center life and threaten equipment availability.*

*Single-module systems are typically able to support availability levels of 99.99 percent, or less than one hour of unplanned downtime annually, and leave critical systems unprotected during UPS maintenance.*

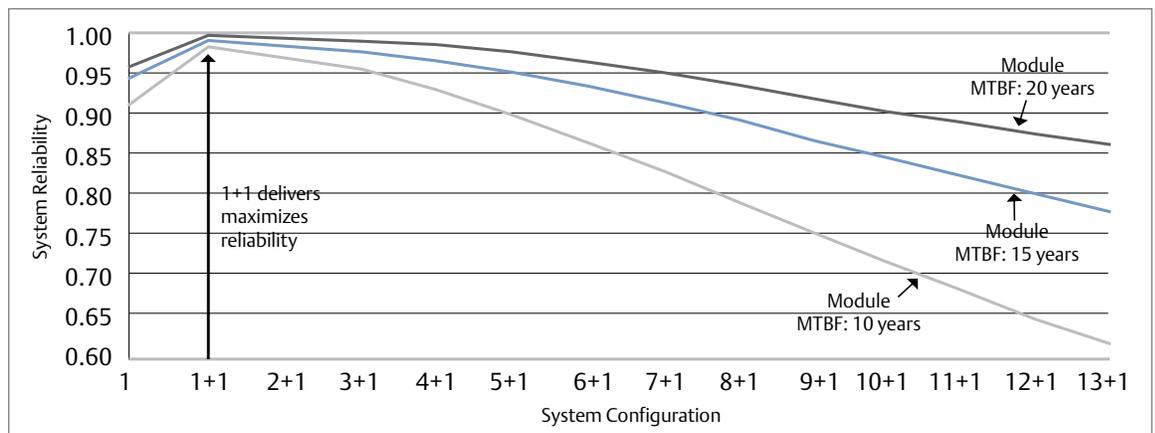
System Type	Description	No. of Buses
SMS	Single module UPS system. No redundancy.	Single
1+1	2 UPS modules running in parallel.	Single
1+N	1 UPS plus N = number of UPS modules added for redundancy. Each module has its own internal bypass. (Accepted practice in Europe and gaining acceptance in US)	Single
N+1	N = number of UPS modules required, plus one additional module for redundancy. In high-availability applications, N is typically less than 3 N + 1 systems use a system-level bypass. (Accepted practice in US)	Single
Dual-bus or 2N or 2 x (N+1)	2 UPS systems feeding 2 independent output distribution systems. UPS output buses are typically in sync. Each bus often contains a redundant UPS system, which is technically a 2 x (N+1), but is sometimes also called 2N. Redundancy can also be accomplished through 1 + N configuration, which is less common but growing in acceptance.	Dual

**Table 1. Summary of most common UPS system configurations.**

keeping parts count and system complexity low. New software-based approaches to scalability allow 1+1 systems to be sized to current requirements while enabling capacity growth up to 100 percent without adding UPS modules.

The N+1 configuration attempts to balance scalability and availability and can effectively accomplish that if system complexity is controlled. This requires proper sizing of the UPS modules. If the initial facility load is light compared to future requirements, it can be tempting to size UPS modules to initial requirements; however this can create reliability problems in the future if the UPS module count gets too high.

Emerson Network Power has analyzed power system reliability and determined that single-bus, multi-module UPS system reliability is acceptable up to a 3+1 configuration (Figure 1). Reliability starts to diminish quickly beyond that point because the increase in the number of modules increases the system parts count, which increases the probability of failure and the risk of maintenance-related failures. Unlike server arrays, which can operate at reduced performance if multiple failures occur, an N+1 UPS system will shut down if UPS module failures put the system below the load capacity point.



**Figure 1. System reliability for N+1 system configurations at 10-, 15- and 20-year MTBF rates. Reliability diminishes quickly beyond the 3 + 1 configuration.**

If reliability is important, UPS modules should be sized no less than one-third of the projected total facility load. If there is high certainty that growth will occur, start with modules sized at one-half the initial load. This provides room for growth while ensuring adequate availability throughout the life of the facility. Batteries, which are one of the most expensive components of the power system, can be sized to initial load with additional capacity added as required.

In addition to reliability considerations, system cost should be factored into decisions regarding UPS module size. The cost per kW for a UPS system drops as module size increases. This makes it more expensive, for example, to implement 10 modules of 10 kW each, rather than one 100 kW modules. Figure 2 shows how cost per kW drops as unit size increases, based on an analysis of actual costs of Liebert UPS products.

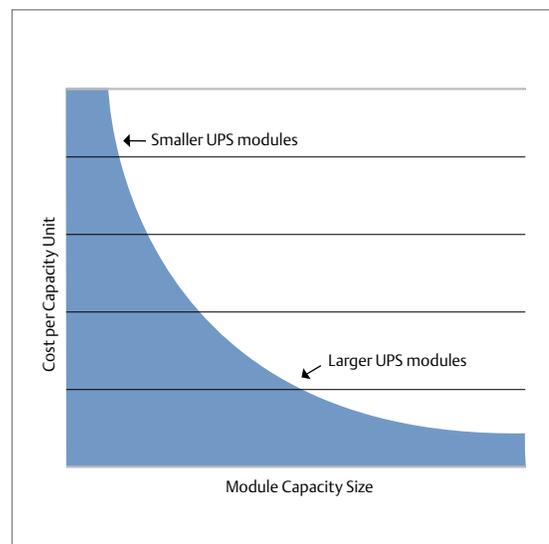
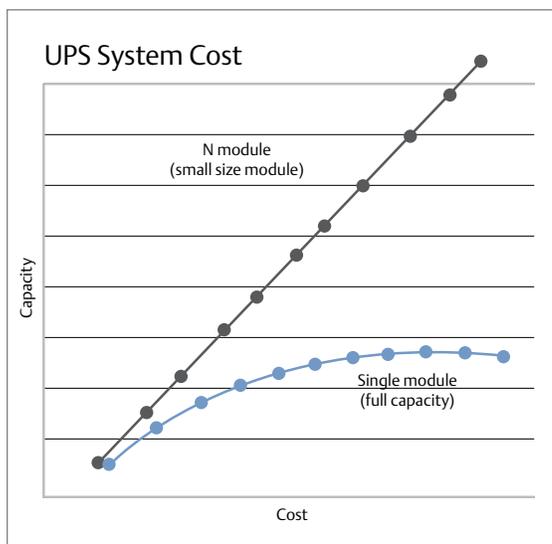
## Compartmentalized Distribution

Traditional power distribution designs use an approach in which the UPS feeds a required number of power distribution units (PDUs), which then distribute power directly to equipment in the rack. This was adequate when the number of servers and racks was relatively low, but with today's equipment it presents scalability and flexibility challenges. Too often breaker space is expended long before system capacity is reached.

Two-stage power distribution represents an emerging alternative that compartmentalizes distribution between the UPS and the server to enable greater flexibility and scalability (Figure 3).

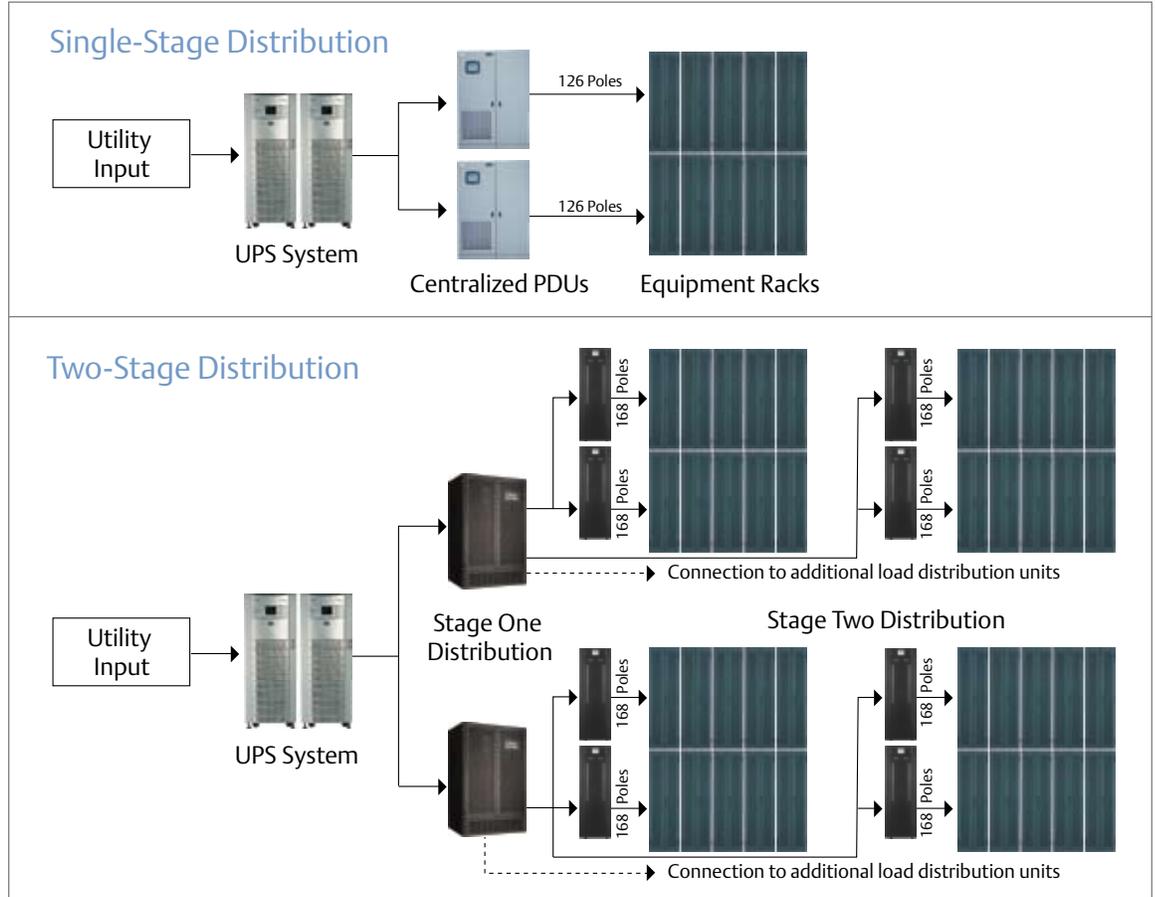
The first stage of the two-stage system provides mid-level distribution. Similar to the traditional PDU, the mid-level distribution unit receives 480 Volt or 600 Volt power from the UPS. The mid-level distribution unit includes most of the components that exist in a traditional PDU, but with an optimized mix of circuit and branch level

*Two-stage power distribution represents an emerging alternative that compartmentalizes distribution between the UPS and the server to enable greater flexibility and scalability.*



**Figure 2. Cost per KW of capacity decreases as UPS module size increases so UPS systems with small building blocks will have a higher total cost than those using optimum size building blocks.**

*The use of two stages separates deliverable capacity and physical distribution capability into separate systems to eliminate the breaker space limitations of the traditional PDU.*



**Figure 3. Compartmentalizing distribution provides a greater number of pole positions for equipment racks while making it easier to adapt the distribution system to changing application requirements.**

distribution breakers. Instead of doing direct load-level distribution, this device feeds floor-mounted distribution cabinets via an I-Line panelboard distribution section. The I-line panelboard provides the flexibility to add up to 10 plug-in output breakers of different ratings as needed.

The load-level distribution unit distributes power directly to the rack-mounted equipment. It is a compact component packaged in a standard rack that can be positioned on the end or in the center of rack rows. When placed in a row, the load-level distribution unit blends in with the equipment racks, enhancing the overall appearance of the facility. It can be tailored to the level of availability the application requires, including single-sourced,

dual-sourced (for dual-bus applications), or fed from four different inputs or sources.

The use of two stages separates deliverable capacity and physical distribution capability into separate systems to eliminate the breaker space limitations of the traditional PDU. The load-level distribution unit can be configured to meet the specific requirements of the technology it directly supports, while the mid-level PDU remains unchanged. For legacy equipment, the load-level unit can be configured with a standard 225A, 42-pole panelboard. For current- and next-generation equipment, it can be configured to support a higher power density, such as 400A or 480V panelboards, while the investment in the mid-level PDU is protected.

The UPS-sizing philosophy covered previously extends to the mid-level PDU. If the mid-level PDUs are sized properly, a series of growth cycles can be supported by adding breakers to the mid-level PDU and additional load-level units as required. This creates the ability to scale the power distribution system without disrupting operations.

Another advantage of this approach is its affect on airflow. With two-stage distribution, under-floor cabling is significantly reduced. Cross-aisle cabling is limited to the feeds that run from the mid-level distribution units to the load-level units. Distribution from the load-level units, which are typically placed alongside IT equipment racks, runs vertically below, or through, the racks, keeping cold aisles free from power cables, as well as reducing the lengths of cable needed to feed the rack-level equipment. This has the added benefit of enhancing the efficiency of the distribution system by enabling fewer, larger power paths from the UPS to the mid-level distribution unit and shorter paths from the load level units to the point of consumption.

### **Rack-Level Power Management**

Changes in rack density create new challenges in rack-level power management. With growth in the overall number of servers and the numbers of servers per rack, racks have become cable management nightmares. That isn't just a headache when changing out equipment—thick bundles of cables can severely limit the ability to move air through the rack, increasing the likelihood of overheating.

In addition, higher densities increase the likelihood that new devices will overload circuits. New compact servers may leave rack space and outlets unused while consuming all available power to the rack. This can present problems when the individuals responsible for deploying new servers are not aware of

power capacity constraints within specific racks. Consequently, servers may be added to existing racks with insufficient capacity, thus overloading the power circuit and tripping the branch circuit breaker, causing the entire rack to go down.

Intelligent power strips present a cost-effective, easy-to-implement solution to in-rack power management (Figure 4). These strips can mount either vertically or horizontally to simplify equipment changes and reduce cable clutter while providing increased visibility and control of rack power consumption. These power strips monitor electrical attributes of an individual power strip, including real-time remote management via SNMP and local LED display of volts, amps, watts and phase loading or the power circuit.

More advanced strips can also provide receptacle-level control of power on/off, enabling receptacles to be monitored and turned on/off locally or remotely to prevent the addition of new devices that could create an overload condition.

*Thick bundles of cables can severely limit the ability to move air through the rack, increasing the likelihood of overheating.*



**Figure 4. Racks with vertically mounted intelligent power strips.**

***Applying monitoring technology at various points in the critical power infrastructure can provide a comprehensive view of the entire data center power infrastructure that supports the data center.***

For some applications, power monitoring may be more effective when implemented at the branch circuit level. This provides less granular visibility than strip-level monitoring, but allows monitoring of more variables. In some high-density environments, a circuit may support only one high density rack; so branch circuit monitoring becomes equivalent to rack-level monitoring.

Power distribution monitoring may be applied to mid-level distribution as well. Most power distribution units provide power monitoring of the main input breaker that provides even more data than branch circuit monitoring. Additionally, monitoring can be extended to sub-feed breakers feeding load-level distribution units. This provides increased visibility and management capabilities that can be used to identify capacity constraints and potential overload situations.

With power requirements continuing to rise, the need for greater downstream visibility into power consumption becomes more pressing. Applying monitoring technology at various points in the critical power infrastructure can provide a comprehensive view of the entire data center power infrastructure that supports the data center. Power monitoring at the branch-circuit or rack level should be considered a critical part of any power infrastructure that is expected to support rapidly changing loads.

## **Conclusions**

Today's data centers are realizing the need for significant investment in order to meet tomorrow's technology changes. Increasing capacity requirements and the need for increased availability have forced data center operations to focus their attention on infrastructure design that allows greater flexibility, higher availability and the lowest total cost of ownership. Best practices in UPS systems configuration are being combined with more robust, flexible power distribution configurations that enable users to better manage their critical infrastructure by deploying the newest monitoring technologies. These strategies can enable both existing and new facilities to support higher densities and capacities while ensuring faster, more cost-effective response to future changes in information technology.

**Emerson Network Power**

1050 Dearborn Drive  
P.O. Box 29186  
Columbus, Ohio 43229  
800.877.9222 (U.S. & Canada Only)  
614.888.0246 (Outside U.S.)  
Fax: 614.841.6022

EmersonNetworkPower.com  
Liebert.com

While every precaution has been taken to ensure accuracy and completeness in this literature, Liebert Corporation assumes no responsibility, and disclaims all liability for damages resulting from use of this information or for any errors or omissions. Specifications subject to change without notice. ©2007 Liebert Corporation. All rights reserved throughout the world. Trademarks or registered trademarks are property of their respective owners. ®Liebert and the Liebert logo are registered trademarks of the Liebert Corporation. Business-Critical Continuity, Emerson Network Power and the Emerson Network Power logo are trademarks and service marks of Emerson Electric Co. ©2007 Emerson Electric Co.

WP152-57

---

**Emerson Network Power.**

The global leader in enabling Business-Critical Continuity™.

- AC Power
- Connectivity
- DC Power

- Embedded Computing
- Embedded Power
- Monitoring

- Outside Plant
- Power Switching & Controls
- Precision Cooling

**EmersonNetworkPower.com**

- Racks & Integrated Cabinets
- Services
- Surge Protection