

Five Technologies Simplifying Infrastructure Management

Achieving high availability is still a priority, but a new generation of infrastructure systems are emerging that enable greater flexibility in how IT systems and facilities are managed.

Introduction

During the last five years, virtually every size business has experienced growth in the number and density of its IT systems as day-to-day operations become more dependent on server, storage and communications technology. As a result, today's data centers must support more devices, consuming more power and generating more heat.

In many cases, large organizations with mature facilities have been forced to deal with these changes by overhauling their infrastructure or developing new facilities sooner than expected.

For smaller, fast-growing businesses, these same issues pose an even greater challenge. Not only must they accommodate the same growth in technology utilization as larger companies, but also plan for technology growth dictated by the growth of the business. In addition, this growth is often occurring within smaller facilities managed by information technology professionals with limited infrastructure experience.

This is risky because the business depends on IT and IT depends on the infrastructure.

“Availability” — ensuring IT systems are available when needed — has traditionally been the dominant driver in infrastructure design. As long as IT systems didn't go down, the infrastructure was doing its job. However, as organizations struggle to address the rapid pace of technology change, a new design consideration has emerged: flexibility.

Achieving high availability is still a priority, but a new generation of infrastructure systems are emerging that enable greater flexibility in how IT systems and facilities are managed. This paper describes five infrastructure technologies that are making it easier for growing businesses to introduce new IT systems as needed while maintaining high levels of availability.

1. Turn On More Power with Software Scalable UPS

Transitioning from rack-level to room-level power protection is a necessary step in the evolution of the IT infrastructure. One of the key decisions that must be made when this occurs is the size of room-level UPS modules to be installed.

Assuming there is adequate power coming into the facility, the power protection system can be scaled by paralleling UPS modules. However, this approach has its limits. If UPS modules are undersized initially, new modules will need to be added frequently, creating a system that is less reliable, more difficult to manage and more expensive. Statistical analysis of UPS system configurations in light of failure rates shows that system reliability begins to decrease sharply when more than 4 UPS modules are used in a single system.

If UPS modules are oversized, the business is paying for capacity it may not need for several years.

Planning for some growth makes sense, but can future growth be accurately predicted? During the Internet boom, many organizations planned for growth that never occurred. The industry then overcompensated at exactly the time equipment densities were rising faster than expected, creating a capacity crunch that is shortening the expected lifespan of some facilities designed in the last five years.

A new approach to UPS scalability has emerged that reduces the risk associated with miscalculating future capacities: software scalability. This solution combines new technologies that flatten the UPS efficiency curve with marketing strategies from UPS manufacturers that are aligned with the needs of growing businesses. Here's how it works.

The UPS module has a potential capacity higher than the purchased capacity, but some capacity is "locked" by the UPS software and can be unlocked as needed with a software key. For example, a growing business could purchase an 80 kVA UPS with only 40 kVA enabled, and then later add capacity in 20 kVA increments up to 80 kVA, effectively doubling the capacity of the UPS without modifying UPS hardware.

The key to this approach is a patented combination of resonant converter technology and unique firmware instructions that allows the UPS to maintain its efficiency at partial loads so efficiency remains high even when some capacity is locked. It also requires that the UPS manufacturer assume the risks associated with predicting future growth requirements, freeing growing businesses to reduce their initial investment without risking the future reliability of their infrastructure system.

A new approach to UPS scalability has emerged that reduces the risk associated with miscalculating future capacities.

Greater flexibility in how modules can be paralleled reduces the risk of being locked into a UPS size that doesn't meet the long-term requirement of the business.

2. Keep Pace with Business Growth with Unlike Paralleling

One factor that has limited some businesses' ability to build a robust, flexible infrastructure is the way UPS modules could be paralleled. Typically, modules could only be paralleled with identical modules. If you installed a 40 kVA UPS initially, you were locked into 40 kVA building blocks that may not prove ideal if capacities grow faster or slower than expected.

Greater flexibility in how modules can be paralleled reduces the risk of being locked into a UPS size that doesn't meet the long-term requirement of the business. Some newer UPS systems now have the built-in intelligence to support this capability.

Together, unlike paralleling and software scalability create a dynamic, controlled approach to growth. Consider a business that installs an 80 kVA software-scalable UPS with 40 kVA enabled. The business can complete two growth cycles (60 kVA and 80 kVA) before adding more hardware. Then another software-scalable 80 kVA UPS with 40 kVA enabled could be paralleled with the initial UPS to bring capacities to 120 kVA. This UPS can also support two additional growth cycles to bring total capacity to 160 kVA.

Unlike paralleling also makes it easier to add redundancy to the UPS system based on current, rather than initial, modules size.



The Liebert NX features both software scalability and unlike paralleling to provide multiple options for adding capacity or availability.

3. Break Free of Circuit Limitations with Compartmentalized Distribution

In 1996, a fully populated server rack could house 14 single-corded servers operating at 120 Volts. That same rack can now house 42 1U servers, which are likely to be dual-corded and operate at 208 Volts, or six blade server chassis. This evolution has created the need for a power infrastructure capable of adjusting to changes in the number, density and location of these devices.

That need is being satisfied by an approach to power distribution that separates primary distribution from load-level distribution.

In traditional power distribution designs, the UPS feeds a required number of power distribution units (PDUs) that 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.


The alternative is to compartmentalize distribution to separate deliverable capacity and physical distribution capability. This is accomplished through a two-stage system in which the first stage receives high-voltage power from the UPS and then feeds floor-mounted distribution cabinets that distribute power directly to the rack-mounted equipment.

The primary distribution unit can be configured with a mix of circuit and branch-level distribution breakers to eliminate breaker space limitations.

Load-level units can be tailored to the level of availability required, including single-sourced or dual-sourced (for dual-bus applications).

A series of growth cycles can be supported by adding breakers to the primary distribution unit and additional load-level units as required. This creates the ability to scale the power distribution system without disrupting operations.

A series of growth cycles can be supported by adding breakers to the primary distribution unit and additional load-level units as required. This creates the ability to scale the power distribution system without disrupting operations.



In compartmentalized distribution, load-level cabinets can be tailored to the requirements of the specific equipment supported.

By effectively cooling high-density racks, supplemental cooling also allows better space utilization, eliminating the need to partially populate racks to spread the heat load.

4. Eliminate Hot Spots and Zones with Supplemental Cooling

Proper environmental control – temperature, humidity and air quality – is critical to the performance and reliability of IT systems. This control has traditionally been provided exclusively through floor- or ceiling-mounted room cooling units.

Often, a raised floor is used to control the distribution of cooling across the facility. This proved sufficient until equipment densities reached current levels and the introduction of new technologies created greater variations in temperatures across the room.

The response to this trend has been the development of supplemental cooling systems that bring cooling closer to the source of heat by positioning small cooling modules either above or alongside equipment racks. This has a number of advantages for a growing business.

First, it allows cooling to be focused where it is needed most. Cooling modules can be placed over every rack in a facility, or only over high-density racks. If rack density changes through equipment moves or additions, cooling modules can be added or relocated to provide greater flexibility and scalability than was possible previously.

By effectively cooling high-density racks, supplemental cooling also allows better space utilization, eliminating the need to partially populate racks to spread the heat load.

Finally, these supplemental cooling systems are more energy efficient because they move air in a local zone rather than in large volumes across the entire data center space and also take advantage of a highly efficient refrigerant fluid to transfer the heat out of the room.



An open approach to supplemental cooling supports can support higher equipment densities across the facility.

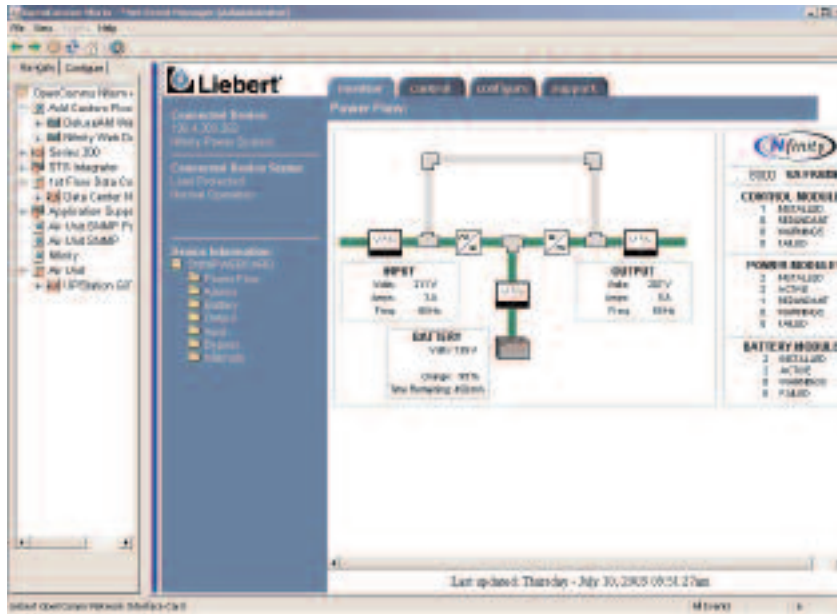
It's important to note that supplemental cooling doesn't replace the need for a room-level system. The supplemental system only removes heat and does not manage humidity or air quality; however, the use of supplemental cooling can greatly reduce the number and capacity of the room systems required.

An open architecture supplemental cooling system uses an overhead piping system to deliver refrigerant to the cooling modules. This allows cooling modules to be positioned anywhere in the room and easily relocated using a quick-connect feature between the piping system and the cooling module.

5. Extend Your Control with Infrastructure Monitoring and Management

Infrastructure systems can be monitored and managed similar to IT systems themselves. Infrastructure monitoring uses the existing network to collect and consolidate alarm and status information from power and cooling systems and deliver this information to any desktop in a format that presents a snapshot view of system health.

The extra protection of a double-conversion UPS is warranted for applications where downtime is especially costly, such as large switches supporting hundreds of phones.



Infrastructure monitoring systems can provide a convenient view of all network-enabled infrastructure systems to simplify management.

Monitoring can also be extended inside equipment racks through the use of intelligent power strips and compact network-enabled environmental monitoring units.

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. The result could be a circuit overload that causes the entire rack to go down.



Intelligent power strips extend power monitoring to the rack.

Intelligent power strips present a cost-effective, easy-to-implement solution to in-rack power management. These strips can mount either vertically or horizontally within the rack to simplify equipment changes and reduce cable clutter while providing increased visibility of rack power consumption. More advanced strips can also provide receptacle-level control of power on/off to enable receptacles to be turned on/off locally or remotely to prevent the addition of new devices that could create an overload condition.

Data from intelligent power strips can be integrated with temperature and humidity readings from in-rack environmental monitors to provide complete visibility of rack operating conditions.

Conclusion

While growing businesses manage a high degree of change with limited resources, a new generation of infrastructure technologies has emerged that make it easier to accommodate unexpected changes in IT systems without disrupting operations. Working together, a software scalable UPS that supports unlike paralleling, compartmentalized distribution, supplemental cooling and infrastructure monitoring can help a growing business create an IT infrastructure that delivers flexibility, high availability and low cost of ownership.

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 mark of the Emerson Electric Co.

Printed in U.S.A. WP153-77

Emerson Network Power.

The global leader in enabling Business-Critical Continuity™.

- | | | | |
|----------------|----------------------|------------------------------|-------------------------------|
| ■ AC Power | ■ Embedded Computing | ■ Outside Plant | ■ Racks & Integrated Cabinets |
| ■ Connectivity | ■ Embedded Power | ■ Power Switching & Controls | ■ Services |
| ■ DC Power | ■ Monitoring | ■ Precision Cooling | ■ Precision Cooling |

EmersonNetworkPower.com