Liebert®

IT White Paper

BALANCING SCALABILITY AND RELIABILITY IN THE CRITICAL POWER SYSTEM: *WHEN DOES 'N + 1' BECOME 'TOO MANY + 1'?*

EMERSON™
Network Power

Summary

Uninterruptible Power Supply (UPS) protection can be delivered through a single-module approach or through redundant systems. Redundancy enables higher availability of critical systems and is typically achieved through either an N + 1 or 1 + 1 design. While 1 + 1 systems deliver a significant improvement in availability over N + 1 systems and are regularly specified for the most critical applications, N + 1 remains a viable and popular option for applications seeking to balance cost, reliability and scalability.

However, the benefits of N + 1 redundancy diminish with the number of modules that are added to the system. In fact, beyond four modules (3 + 1), the complexity of an N + 1 system begins to significantly compromise reliability. System costs and service requirements also increase with the number of UPS and battery modules added. Increased service requirements typically mean increased human intervention, increasing the risk of downtime.

When N + 1 redundancy is used, UPS modules should be sized so that the total anticipated load can be carried by three modules.

Consequently, when N + 1 redundancy is used, UPS modules should be sized so that the total anticipated load can be carried by at most three modules. While UPS systems are technically scalable beyond this point, 3 + 1 should be considered the threshold at which scalability has such a negative impact on system availability, cost and performance that it is not recommended.

For example, in a data center or computer room that is expected to eventually support 600 kW of critical load, the switchgear and distribution panels are sized to this anticipated load. The UPS configuration that maximizes availability for this room is two 600 kW UPSs operating in parallel (1+ 1). If budget limitations do not allow this configuration, an N + 1 configuration should be considered to effectively balance cost, reliability and scalability. In this case, initial capacities could be met by two 200 kW UPS modules operating in a 1+ 1 configuration. As the load in the room increases, additional 200 kW UPS modules can be added until the capacity of the room is reached. At full capacity the UPS system will include four UPS modules operating in a 3 + 1 configuration. This achieves a balance between reliability and cost management in the context of unpredictable data center growth.

UPS Redundancy Options

All UPS equipment and switchgear, regardless of manufacturer, requires regular preventive maintenance during which the UPS system must be taken offline. Redundancy enables individual UPS modules to be taken offline for service without affecting the quality of power delivered to connected equipment. Redundancy also adds fault tolerance to the UPS system, enabling the system to survive a failure of any UPS module without affecting power quality.

There are a number of different approaches to UPS redundancy, which are summarized in Table 1. For more information on these options, see the Emerson Network Power white paper, *The Role of Redundancy in Increasing Power System Availability*, which describes each approach.

This paper provides a detailed analysis of the commonly used parallel redundant option (N + 1), focusing on the reliability, cost-effectiveness and service requirements of this architecture. The analysis is based on the availability of conditioned power and therefore bypass is not considered for any option.

In a parallel redundant (N + 1) system, multiple UPS modules are sized so that there are enough modules to power connected equipment (N), plus one additional module for redundancy (+ 1). During normal operation the load is shared equally across all modules. If a single module fails or needs to be taken offline for service, the system can continue to provide an adequate supply of conditioned power to connected systems.

N + 1 systems can be configured with either a scalable or modular architecture. The scalable architecture features UPS modules that each include a controller, rectifier, inverter and battery. In the modular architecture, the power modules comprise

System Configuration	Concurrent Maintenance?			Fault Tolerance?		Availability
	Module	**System**	**Distribution**	**Module**	**Distribution**	
Single module	No	No	No	High	No	high
Parallel redundant	Yes	Yes	No	Yes	No	higher
Small isolated redundant	Yes	Yes	No	Yes	No	higher
Large isolated redundant	Yes	Yes	No	Yes	No	higher
Distributed redundant	Yes	Yes	Yes	Yes	Yes	continuous
Selective redundant	Yes	Some	Selective	Some	Selective	continuous
Power-Tie™	Yes	Yes	Yes	Yes	Yes	continuous
Hybrid AC-DC Power System	Yes	Yes	Yes	Yes	Yes	continuous

Table 1. Summary of system configurations.

a rectifier and inverter. A single or redundant controller controls operation and the battery system is shared among modules. All modules in an N + 1 system share a common distribution system.

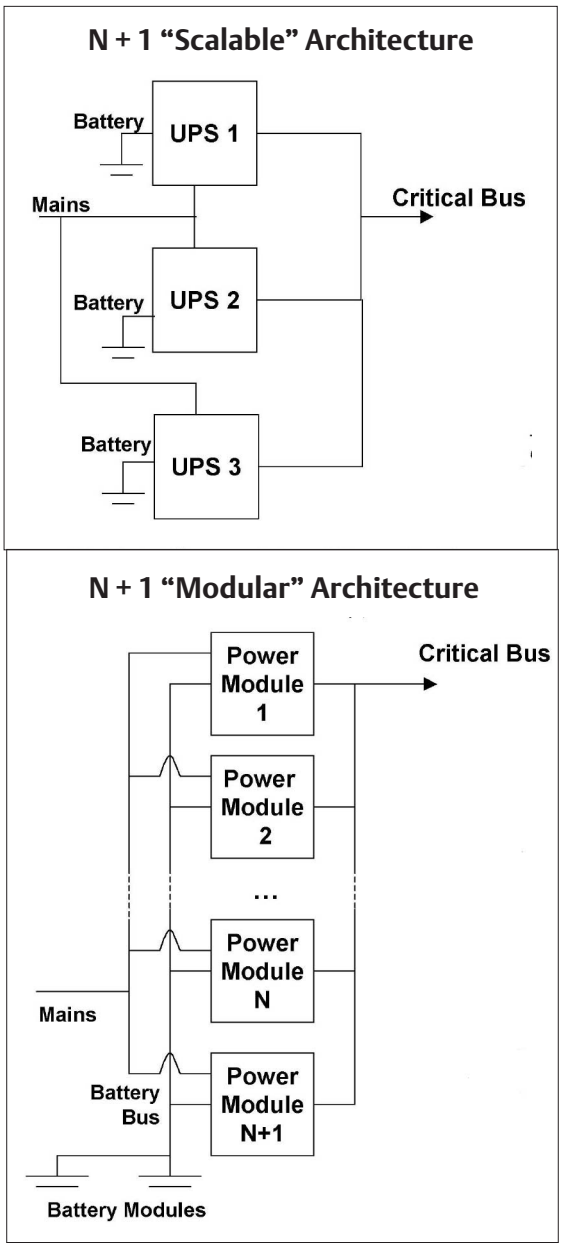

Figure 1. *Top:* In a scalable N + 1 architecture each UPS has its own controller and battery systems. *Bottom:* In a modular N + 1 architecture, power modules may share a controller and battery system. The common battery bank constitutes a single point of failure.

Scalability: IT Systems vs Power Systems

Network and data center managers expect scalability in the systems they use to manage and route data because future requirements are dynamic and difficult to project. A system that can easily "scale" to meet increased requirements enables an organization to invest in technology based on current needs without having to abandon or replace those systems when requirements change.

This is obviously desirable and is often promoted as a benefit of the N + 1 architecture. With this architecture, the argument goes, the UPS system can be sized to initial capacities and additional UPSs or power modules can be added later as capacities increase. This is true to a point. To find that point it is first necessary to understand the difference in how "scalability" applies to IT systems versus power systems.

For IT systems, scalability refers to the ability to add processors, memory, or controllers without swapping out the rack or enclosure. In the power infrastructure, it refers to the ability to add power and battery modules as the load increases. While similar conceptually, there are significant differences in how scalability applies to power systems and IT systems.

- **Failure Consequences**. In IT systems, the failure of a single module may create a slight degradation in the performance of the system. A failure of a second module increases this degradation and may make the application unavailable to some users.

While similar conceptually, there are significant differences in how scalability applies to power systems and IT systems.

The investment in support systems needs to be weighed against the value of those systems to the business.

In an N + 1 power system, a failure of one UPS module has no effect on system performance; however, a failure of two modules results in a *complete shutdown of all systems that depend on the UPS.* The N + 1 system will not support any equipment if two modules fail, regardless of whether the system has two or fifteen modules.

- **Open vs Closed Scalability**: In IT hardware systems, standardization often enables additional memory or processor modules to be added from a manufacturer other than the original equipment manufacturer. In the power system, additional modules must be acquired from the original manufacturer and must be for the same model UPS.
- **Expected Lifespan:** IT systems are typically upgraded every three to five years, while the power infrastructure must serve the entire life of the data center, often ten years or more. This makes the previous point even more significant. Will equipment manufacturers support modules for 10 years or more? Will expansion modules be available and at what cost? Will vendors guarantee backward compatibility for that period?
- **Software Cost Optimization**: Software licensing costs are becoming an increasingly large component of IT budgets. IT managers need incrementally scalable computing hardware to optimize costs of software licenses that are charged on the basis of number of CPUs or MIPS. There is no such issue with the power system.
- **Expansion Capability:** While it can be difficult to project future capacities, it is a necessary step in the design of a data center or computer room. Future capacities are projected and used to size the main input switchgear and main power distribution panels in data centers and server rooms. The UPS system cannot expand beyond the capacity of these components.

These factors, taken together, make scalability completely different for the infrastructure than for the IT systems the infrastructure supports. Certainly scalability is a desirable trait, but it is desirable only if it can be achieved without compromising availability.

Network Criticality and the Cost of Downtime

Organizations don't acquire a UPS system for its own sake; they acquire a UPS system because they understand the importance of ensuring power quality and availability to the hardware and software systems that are critical to business operations. The more important these systems are to the business, the more important the investment in support systems to overall business success.

As a result, this investment needs to be weighed against the value of these systems to the business. That value comes in two forms. First, and most obviously, is avoidance of downtime and the associated costs, which include loss of productivity, loss of revenue and loss of customer confidence. These costs can be extremely high for businesses that are moderately to

highly dependent on network systems and provide a basis for making sound business decisions in relation to the network and support system. However, studies indicate a surprising number of organizations do not accurately quantify the cost of network downtime. A recent survey by Forrester Research revealed that 67 percent of enterprises either did not know or could not provide an estimate of the costs of downtime to the business.

Not only is it important to analyze downtime costs, these costs should be considered relative to overall business costs. Network criticality is not necessarily a function of the size of the data center or computer room. It is a measure of cost of downtime versus expected profits/gains. A small computer room can be just as critical to the business it supports as a large data center.

The second value of the support system is enabling organization to do more with technology. Maintaining 7x24 availability of network services and deploying new business applications such as IP telephony are only possible through an appropriate support system.

Availability and the Power System Infrastructure

The relationship between IT system availability and the power system infrastructure is illustrated in Figure 2. The desired business result is at the top of the pyramid, the application layer is in the middle and the critical support infrastructure is at the bottom. The pyramid is inverted because the investment is smallest at the bottom of the pyramid and largest at the top.

Interestingly, the relative costs for each layer of the pyramid tend to remain fairly constant regardless of the size of the facility. This is significant because every size data center is now being asked to support higher levels of availability and a misconception persists that it is relatively more expensive for smaller data centers to achieve higher levels of availability than for larger centers. Typically the proportion of capital expenditures dedicated to critical power systems is 2 to 4 percent of the total capital expenditure in the data center, regardless of the size of the facility.

The relative costs for each layer of the Pyramid tend to remain fairly constant regardless of the size of the facility.

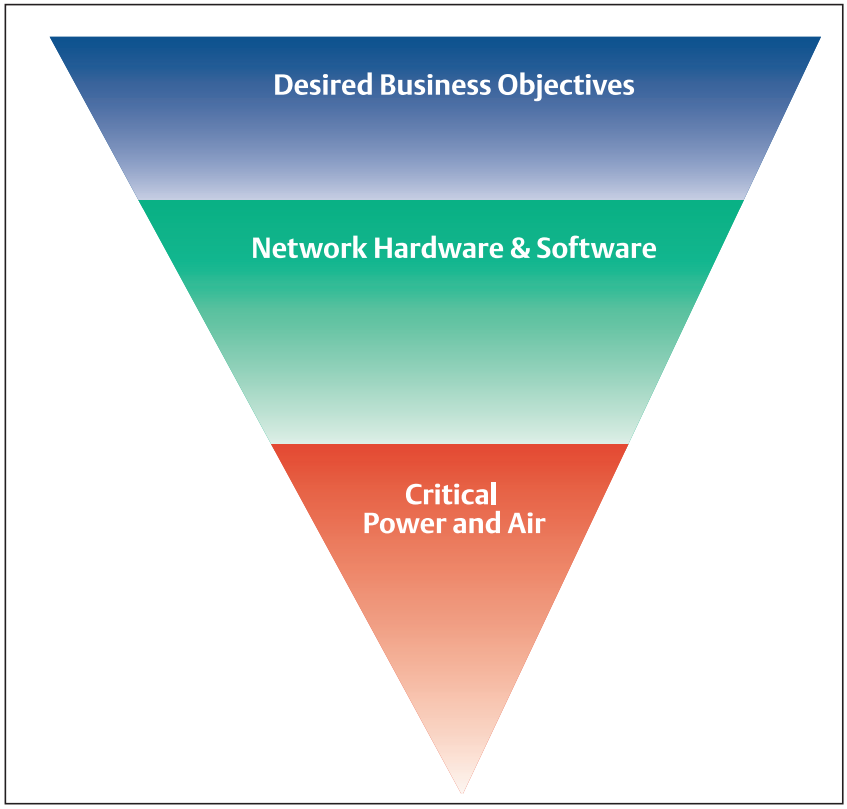

Figure 2. Critical support systems represent a much smaller investment than the network application layer, but are the foundation that supports the application layer's ability to achieve business objectives.

Critical power system availability must be 100 times greater than the availability of the systems being supported to keep from negatively impacting total system availability.

This also puts into perspective the cost of power system "future sizing"– the practice of sizing the power system based on projected capacities rather than capacities required at startup. This *may* add up to 1 percent to total data center capital expenditures – definitely worth saving if possible. But only if this can be accomplished without compromising the availability of the middle of the pyramid. As will be seen in the following sections, a power system that does not adequately consider future growth will compromise overall availability – and ultimately cost more than a system that is properly sized.

System availability is calculated by dividing the hours in a year the system is available by the total number of hours in a year. Because availability of the systems in the middle of the pyramid is dependent on the systems at the bottom, the availability of network hardware and software is the product of the availability of those systems multiplied by the availability of the critical power system. This relationship is illustrated in Table 2. Critical power system availability must be 100 times greater than the availability of the systems being supported to keep from negatively impacting total system availability.

Calculating Availability of the N + 1 Architecture

In terms of the power infrastructure, availability can be projected based on the system design and the reliability of system components. Reliability is measured in terms of Mean Time Between Failure and Mean Time to Repair. Availability is also calculated as follows:

$$\frac{MTBF - MTTR}{MTBF}$$

The critical bus is available if at most one power system is down. The probability of this is equal to the probability that each power system is up, plus the probability that one power system is down. If R is the probability of single UPS plus battery availability, the availability of a 1 + 1 system will be

$$R^2 + 2 \times R \times (1 - R)$$

And the availability of a 3 + 1 system will be

$$R^4 + 4 \times R^3 \times (1 - R)$$

Figure 3 shows how this translates into power system availability for N +1 systems from 1 + 1 to 13 + 1.

Critical Power Availability	IT System Availability	Total Availability
.99	.9999	.9899
.999	.9999	.9989
.9999	.9999	.9998
.99999	.9999	.99989
.999999	.9999	.9999

Table 2. Total availability of IT systems is a product of the availability of the network hardware and software multiplied by the availability of critical power systems.

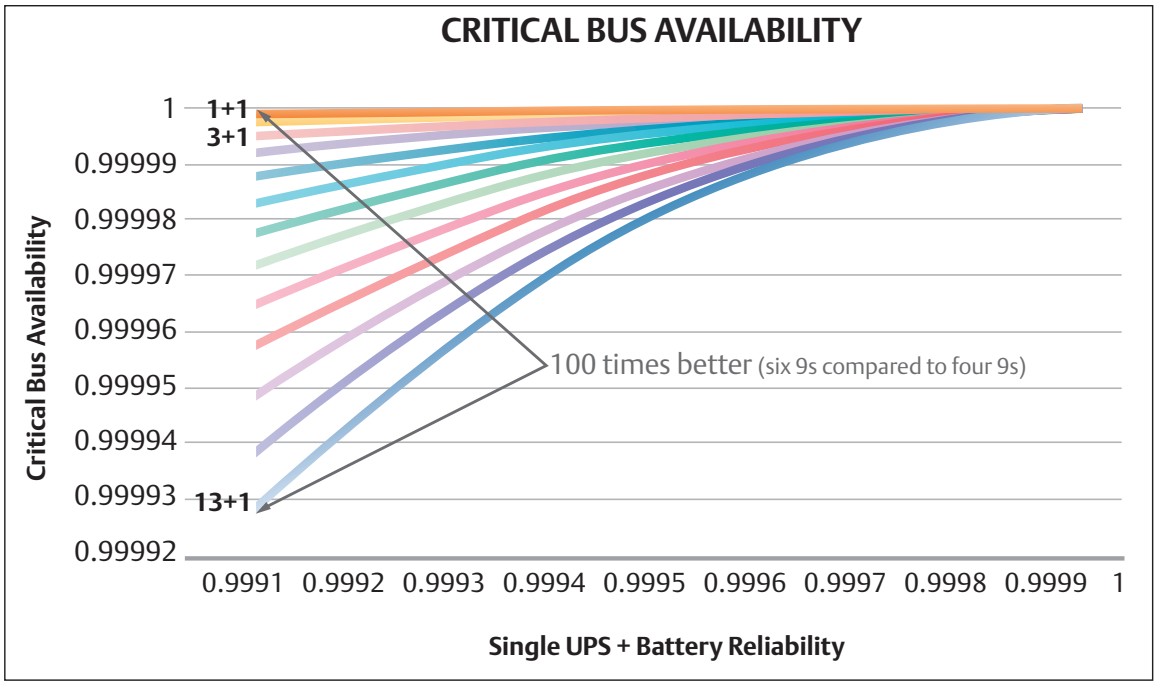

Figure 3. Critical bus availability drops as more modules are added to an N + 1 system. Beyond 3 + 1, the drop in availability begins to represent a significant risk for the business.

At 4 + 1, power system availability begins dropping precipitously. At 13 + 1, power system availability is four nines as opposed to six nines for a 1 + 1 system.

Critical bus availability drops as the number of modules goes up; however, the curve stays fairly flat up to the 3 + 1 level. At 4 + 1 critical bus availability begins dropping precipitously. At 13 + 1, critical bus availability is four nines as opposed to six nines for a 1 + 1 system (assuming a single UPS plus battery system reliability of 3 nines).

This is particularly problematic because modules are added to an N + 1 system as the load increases. Typically an increase in load correlates with an increase in network criticality (i.e. cost of downtime increases). So, an N + 1 architecture is responding to an increase in network criticality by reducing critical power bus availability.

If the reliability of a single UPS and battery system is .9995, a 13 + 1 system will be down about 90 times more than a 1 + 1 system (see Figure 4).

Calculating the Cost of the N + 1 Architecture

Even with the reduced availability of 4 +1 and higher modular systems, some organizations might be willing to risk a pay-as-you-grow approach to power system design if significant costs savings could be realized. However, it isn't just availability that drops as the number of modules increases; cost-effectiveness goes down as well. This is because UPS costs go down on a per-kW basis as the size of the UPS increases. This is also true for battery systems: cost per ampere/hour goes down as ampere/hour rating goes up. As a result, the cost of protecting a 500 kW room may well be less for a 1 + 1 system using two 500 kW UPS plus battery systems than if 14 units of 40 kW UPS, along with 14 battery modules, are used in a 13 + 1 architecture.

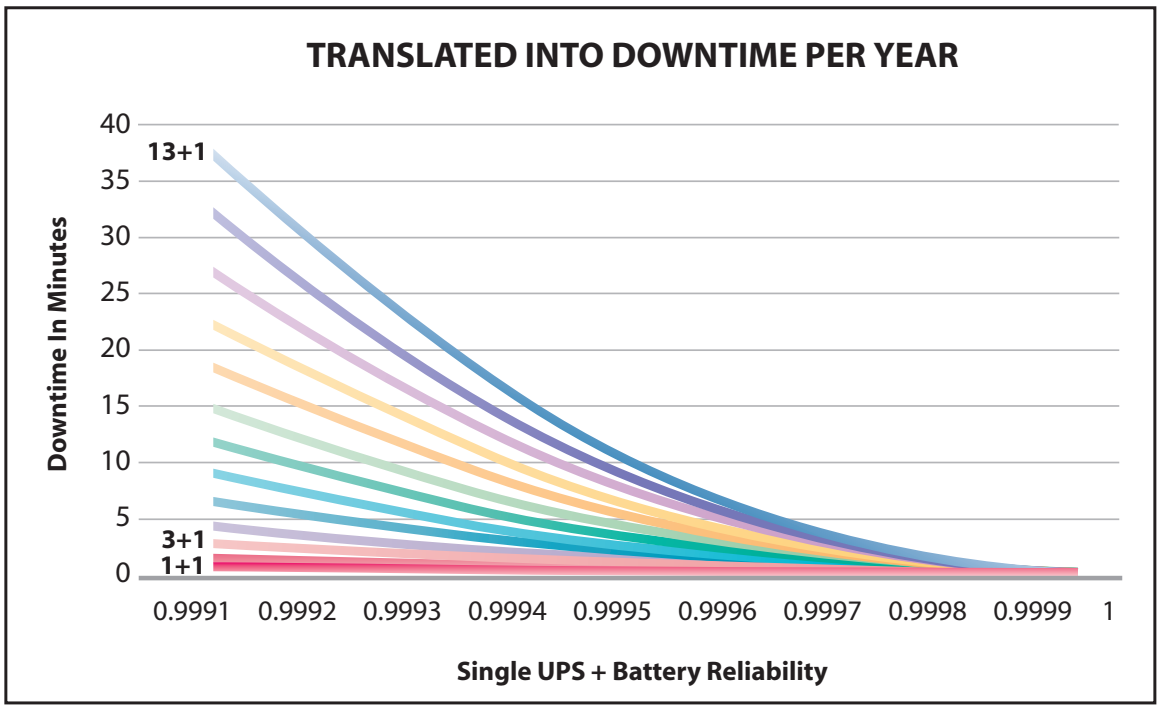

Figure 4. Differences in critical bus availability based on the number of UPS modules creates significant differences in the amount of downtime that can be expected.

The Final Nail: Service Requirements

In studies of the causes of critical system downtime, human error typically represents from 15 to 20 percent of occurrences, behind only hardware and software failures. Unfortunately, the N + 1 architecture increases the likelihood of human error-related downtime – the more modules in a system, the higher the probability of human error due to increased human intervention for service.

This can also be analyzed statistically. If R is the probability of single UPS plus battery availability, service will be required whenever any unit goes down. For a 4 + 1 system, this can be calculated as follows:

$$1 - R^5$$

Performing this calculation on various N + 1 configurations produces the graph in Figure 5. This graph shows that a 13 + 1 architecture is 6.6 times more likely to require service attention than a 1 + 1 system and 3.3 times more likely than a 3 + 1 system. Also, remember that Figure 5 does not factor in the increased probability of downtime resulting from other activities, such as the addition of new power or battery modules.

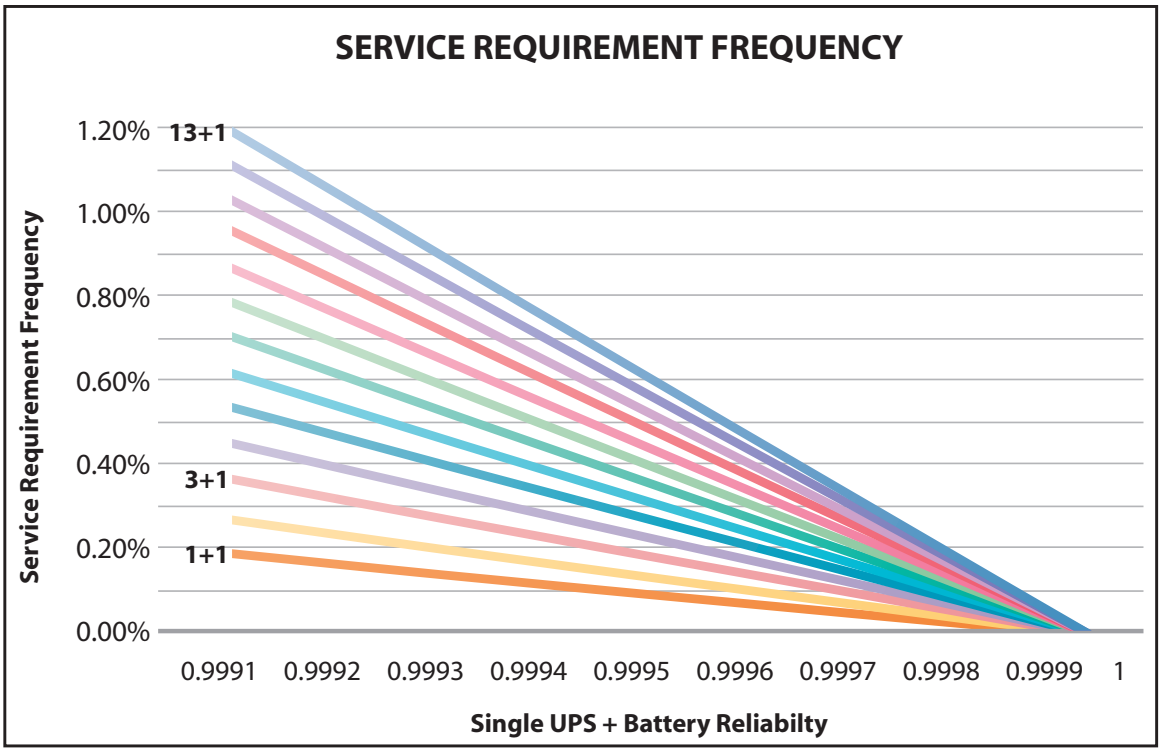

Figure 5. Service requirements also increase with the number of modules, increasing the possibility of downtime from human error.

Conclusion

Infrastructure costs are a relatively small percentage of total data center capital expenditures. But, they have a significant impact on IT system utilization and availability and therefore on the business itself.

Organizations should seek ways to minimize infrastructure costs where possible, but only if this can be accomplished without compromising availability. Decisions that reduce critical system availability may end up reducing the return on investment in all IT systems and limiting the ability to achieve business objectives.

Availability of the N + 1 architecture becomes unacceptable at system configurations that utilize five or more modules. These configurations present greater risk of downtime, are less cost-effective and require more service attention than systems that use four or fewer modules. As a result, the recommended design standard is to use a 1 + 1 configuration whenever possible. If initial capital limitations dictate an N + 1 architecture, UPS modules should be sized so that no more than three modules are required to support the expected capacity of the room. If a 1 + 1 configuration is used to meet initial requirements, the system can accommodate growth of 300 percent without significantly compromising availability.

Liebert Corporation
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
www.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

© 2004 Liebert Corporation. All rights reserved throughout the world. Specifications subject to change without notice.

Trademarks or registered trademarks are property of their respective owners.

® Liebert and the Liebert logo are registered trademarks of the Liebert Corporation

® Keeping Business in Business is a registered trademark of the Liebert Corporation.

The Emerson logo is a trademark and service mark of the Emerson Electric Co.

Printed in U.S.A. 0904 WP403