Cost, Speed, and Reliability Tradeoffs between N+1 UPS Configurations

White Paper 234

Revision 0

by Wendy Torell

Executive summary

There is an increasing trend towards N+1 UPS architectures – rather than 2N – as IT fault tolerance through software continues to improve. There are two common ways N+1 can be achieved: paralleling multiple unitary UPSs together *or* deploying a single UPS frame with multiple internal modules configured for N+1 redundancy. In this paper, we quantify key tradeoffs between an internal "modular" redundant UPS and parallel redundant UPSs, and show a 27% capital cost savings and a 1-2 week decrease in deployment time when internal redundancy is deployed. We also discuss the importance of fault tolerance within the UPS to ensure availability, reliability, and maintainability needs are met.

Introduction

Data centers implement varying degrees of redundancy, based on the criticality of the load(s) they support. Dual path architectures (two separate power paths, for example) provide the highest levels of availability, as maintenance or failure can occur with any system without taking down the load.

In today's data centers, however, we are seeing more and more fault tolerance occur through software, at the IT layer. With technologies such as virtualization and hyper-convergence, no longer is it true that a server going down means the IT mission goes down. If a physical server goes down due to an upstream failure, or is scheduled to go down for maintenance, the data center is able to migrate the business function(s) over to another server, another pod, another room, or a completely separate data center.

While availability is still the critical objective of data centers, some are finding that this can now be achieved with N+1 redundancy of key physical infrastructure systems like the UPS. In this paper, we will clarify the different methods of achieving N+1 redundancy of your UPS system(s), quantify the capital cost, deployment time, efficiency, and reliability tradeoffs, and discuss the importance of fault tolerance within the UPS to ensure reliability, availability, and maintainability needs are met.

Terminology clarification

In many data center discussions, the term "N+1" is used interchangeably with various UPS configurations. Below we define key terms to clarify the distinction between three specific "N+1" configurations.

Battery design

Regardless of the specific N+1 configuration chosen, the battery system can be deployed in a way that improves the fault tolerance.

Today it is common practice to have multiple battery strings, cabled in parallel to provide added redundancy for maintenance.

With this approach, if a battery string fails, the remaining strings can continue to support the load in the event of an outage. Note, the runtime would be decreased with a failed string, but can provide sufficient time so downtime is avoided.

- **N+1 redundant:** a means of achieving resilience that ensures system availability in the event of component failure. Components (N) have at least one independent backup component (+1).[1] Simply speaking, N refers to my *need*, and +1 means I have *one spare*.

- **Isolated redundant:** a specific N+1 configuration. In this configuration, there is a main or "primary" UPS module that normally feeds the load. This configuration requires that the primary UPS module have a separate input for the static bypass circuit. The "isolation" or "secondary" UPS feeds the static bypass of the main UPS module(s) and is completely unloaded.[2]

- **Parallel redundant:** a specific N+1 configuration. It consists of paralleling multiple, same size UPS modules onto a common output bus. Parallel redundant systems require UPS modules identical in capacity and model.[2]

- **Internally "modular" redundant:** a specific N+1 configuration. This is a new term we are defining here because of lack of common nomenclature; in this configuration, the "+1" occurs inside the UPS frame, generally at the power module level. In this configuration, there is a shared backplane, control system, and battery plant.

---

[1] https://en.wikipedia.org/wiki/N%2B1_redundancy

[2] White Paper 75, Comparing UPS System Design Configurations, discusses isolated redundant and parallel redundant configurations in greater detail.

In this paper, we focus on the comparison of **parallel redundant** and **internally "modular" redundant** – the two N+1 configurations where the "+1" components are active (as opposed to standby). **Figure 1** illustrates conceptually where the "+1" occurs in each case.

Parallel redundant   Internal "modular" redundant
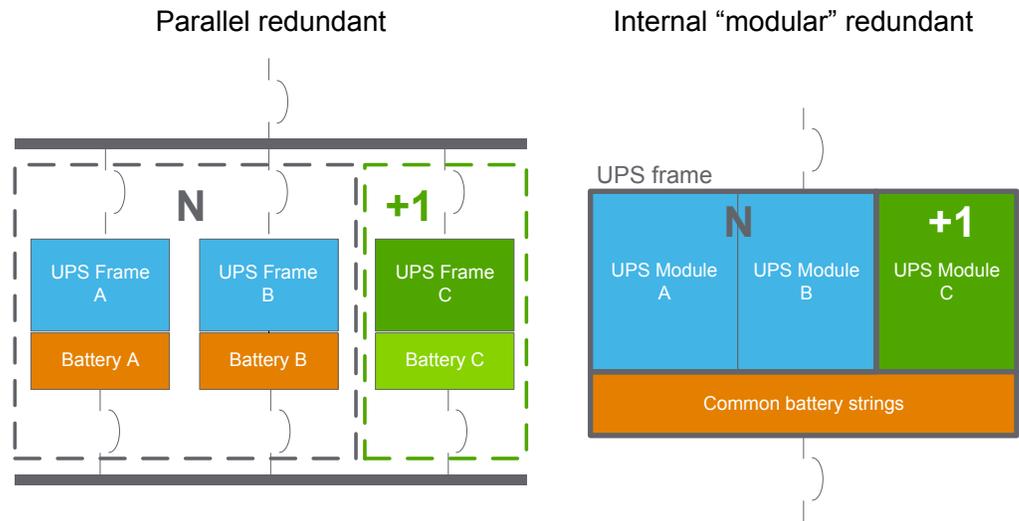
Figure 1

Conceptual difference of parallel redundant and internally "modular" redundant

Configuration details

In order to highlight the trade-offs when choosing a UPS with no redundancy vs. an N+1 UPS, we've analyzed three specific configurations. In all three cases, we chose 1MW for the rated capacity[3].

1. **Baseline 1N configuration:** A single 1000 kW UPS with no redundancy (made up of four "internal" 250 kW modules); the baseline case
2. **Internally "modular" redundant N+1 configuration:** A modular 1000 kW UPS, consisting of five "internal" 250 kW modules (four for capacity and one for redundancy)
3. **Parallel redundant N+1 configuration:** Three 500 kW UPS "frames" configured as parallel redundant (two for capacity, one for redundancy)

Baseline 1N configuration

With a 1N UPS design, failure of any component requires shifting the load to UPS bypass or wrap-around bypass. The simplest case is a single UPS rated to support the entire load. Some UPSs are designed to be modular and scale, while others are of fixed capacity. The 1N UPS we analyzed was a modular design, consisting of four 250 kW modules inside a single frame to achieve the 1000 kW of rated capacity. 1N can also be achieved by paralleling multiple units together to collectively achieve the capacity required. **Figure 2** illustrates the 1N UPS we analyzed.

---

[3] Galaxy VX family was used for estimated street pricing.

Life Is On   Schneider Electric

Figure 2

Baseline "N" configuration
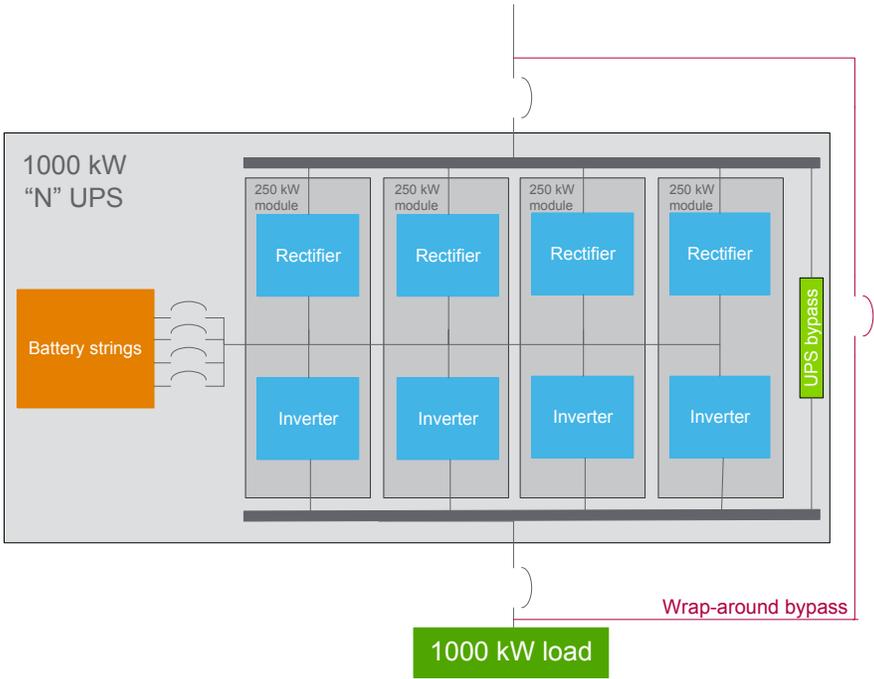
## Internally "modular" redundant N+1 configuration

Modular UPSs, like the one described in the baseline case, can also provide N+1 redundancy, by adding an additional power module. See **Figure 3**. This internally modular redundant N+1 configuration is the same as the baseline case with the addition of a fifth 250 kW module. The battery system consists of four paralleled battery strings, each with their own breaker, so that failure in one does not bring the entire battery system down. As the figure shows, the UPS output bus and the battery (DC) bus is common to all of the modules, therefore they represent single points of failure to the UPS system that would require shifting to wrap-around bypass.

Figure 3

Internally "modular" redundant configuration

Life Is On   Schneider Electric
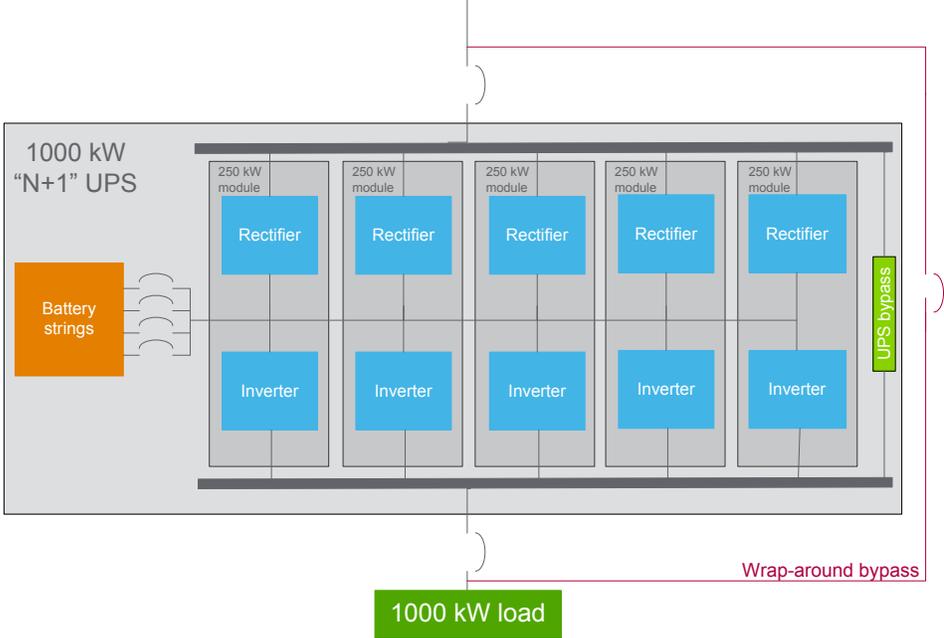
Parallel redundant N+1 configuration

Figure 4 illustrates the third case we analyzed. In this case, three separate UPSs are paralleled together, with a common output bus. Each UPS is 500 kW in capacity, so the third UPS is there for "+1" redundancy. As the figure illustrates, with a parallel redundant configuration, each UPS has its own battery system (each consisting of three paralleled strings), which provides an additional level of redundancy that doesn't exist with the modular UPS. The UPS output bus still represents a single point of failure, as in the other two configurations, except in this case, the output bus is external to the UPS and field-installed.
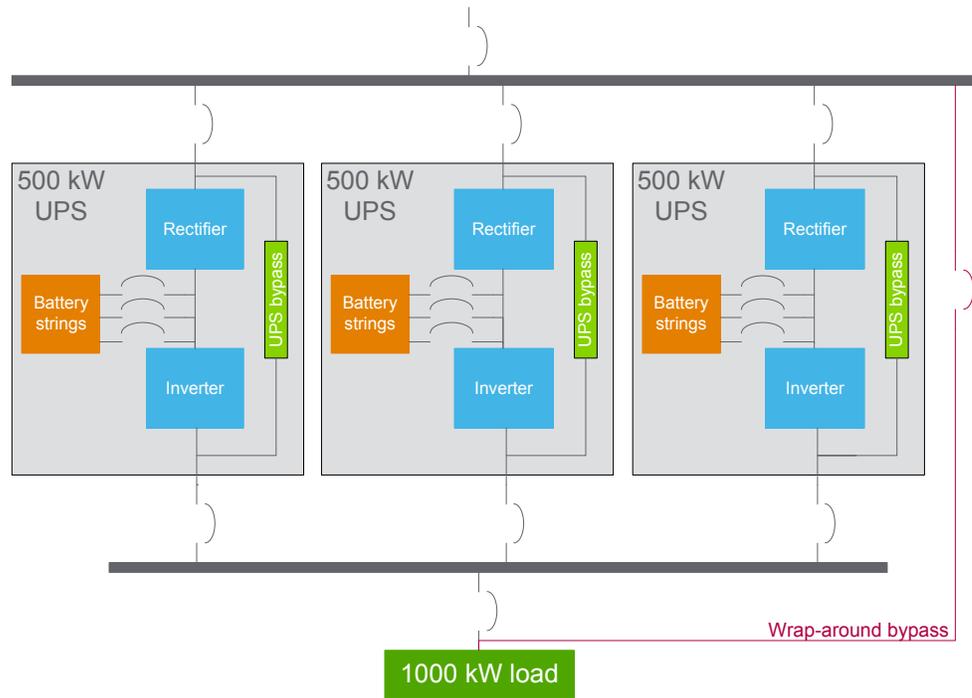
Figure 4

Parallel redundant configuration

Note, sometimes parallel redundant UPSs are deployed with a common battery bank. The advantage of doing that is the cost savings (less battery expense), however, the configuration is now more similar to the internal "modular" configuration in terms of fault tolerance / reliability. The analysis that follows assumes each UPS has its own battery system.

Capex comparison

In general, the more redundancy built into a UPS configuration, the more costly it will be. But it can be challenging for data center managers to make the business case for a particular level of redundancy. In this section, we provide a capital cost analysis of the three configurations discussed earlier, to help decision makers weigh the cost / benefit tradeoffs.

Methodology and assumptions

Detailed one-line diagrams of the configurations were used when estimating the costs of each design. The capital cost includes materials and installation. Installation cost includes labor and all cabling, conduits, hangers, lugs, etc. Material cost includes UPS(s), maintenance bypass cabinet, output breaker, battery system, and assembly service. We have excluded the input breaker, as it is generally assumed to be building premise.

Life Is On   Schneider Electric

Additional costs that were not included in this analysis are: rigging, storage, ongoing maintenance, and space. Even with an ideal layout for each configuration, configuration 3 will need about 25% more space than the other two configurations, which represents additional savings for the data center when building out the space.

The key assumptions of the analysis are:

- Installation cost based on average US electrical install rate in Northern Virginia.
- All wires placed in electrical metallic tubing (EMT).
- UPS input switchgear located 15 meters (50 feet) away from main switchgear. 1600A 3 wire + ground for configurations 1 & 2, and 2000A 3 wire + ground for configuration 3.
- UPS located 3 meters (10 feet) away from UPS input switchgear. 1600A 3 wire + ground to the 1000 kW UPS for configurations 1 & 2, and 800A 3 wire + ground to each of the three 500 kW UPSs for configuration 3.
- Output UPS switchgear located 3 meters (10 feet) away from UPS. 1600A 3 wire + ground for configurations 1 & 2, and 700A 3 wire + ground for configuration 3.
- Wrap-around (maintenance) bypass distance of 6 meters (20 feet). 1600A 3 wire + ground for configurations 1 & 2, and 2000A 3 wire + ground for configuration 3.
- Load located 15 meters (50 feet) away from UPS output switchgear. 1600A 3 wire + ground for configurations 1 & 2; 2000A 3 wire + ground for configuration 3.

Findings

Figure 5 summarizes the differences between the capital costs of the three configurations. As the chart illustrates, internal "modular" redundancy is $128/kW (26.9%) lower in capital cost than the parallel redundant configuration, and the baseline case is $29/kW (6.1%) lower in capital cost than internal "modular" redundancy.
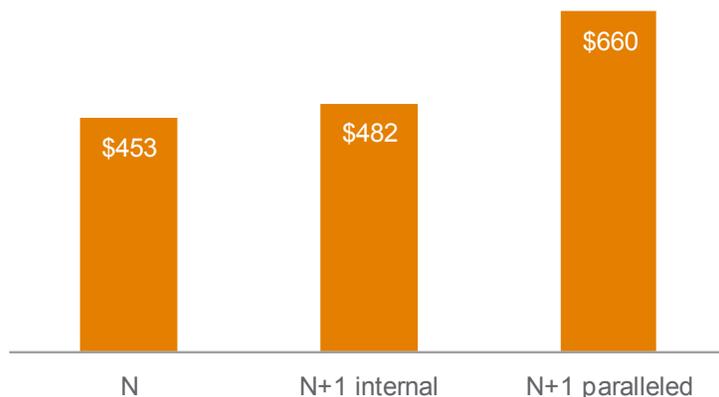
Figure 5

Cost / kilowatt comparison of 3 configurations

Table 1 provides a further breakdown of the estimated costs for each design by major cost category. All costs were normalized to cost/kilowatt of rated UPS capacity. While there is a cost per kW difference between small and large capacity UPSs, this provides a reasonable guideline for the relative cost difference between approaches.

Table 1
Detailed results of cost comparison

Cost categories	1N UPS	N+1 internal "modular"	N+1 parallel redundant
UPS	$148/kW	$177/kW	$214/kW
Bypass cabinet(s)	$36/kW	$36/kW	$70/kW
7 minute battery system	$98/kW	$98/kW	$148/kW
Assembly service	$8/kW	$9/kW	$16/kW
Installation cost	$163/kW	$163/kW	$213/kW

Speed of deployment

In addition to capital cost differences between N+1 configurations, there are also implications to the speed of deployment. In this section, we discuss what goes into installing a single UPS as compared to a set of parallel redundant UPSs.

A typical installation of a 1 MW UPS requires a span of approximately 6-8 weeks (which includes buffers between critical steps). The main activities that occur during this timeframe include:

- **Prep of the room** for the UPS system, including the concrete housekeeping pad(s). A week on the project schedule is generally allotted for this activity. Then there is typically a buffer of a week between room prep and delivery to ensure the unit doesn't show up to a room that isn't ready.
- **Delivery and rigging** of the UPS. A 1000 kW UPS system is very heavy and cumbersome. 2 to 3 days are generally allotted on the project schedule for this step.
- **Running conduits for the UPS.** The feeders will take 3 sets of 3" conduits for each feed. This work can take on the order of a week.
- **Pulling the wires and making the terminations.** A week is generally allotted on the schedule for this work.
- **Schedule startup and testing;** Project schedules generally incorporate a buffer of a week between completely connecting the UPSs and scheduling the startup. This is to account for any unforeseen issues that may arise during installation. Testing then takes a week.

These installation steps are the same for the 1N design and the internal "modular" redundant UPS, with the exception of adding an extra power module in the frame. Therefore, the installation costs are the same. For a parallel redundant UPS configuration, where large UPSs must be paralleled together, the typical deployment time increase is on the order of 1 – 2 weeks or **25%-30% longer for a paralleled system**. The additional field work that goes into the multi-unit installation to setup, configure, and ensure communication between units are below:

- More terminations for more electrical feeds
- More units to set in place
- More units to startup
- More units to load bank test
- Paralleling and sync checks
- More procedures to test/perform
- More control wiring and monitoring points.

Life Is On | Schneider Electric

With a modular UPS, where multiple internal "modules" are used to increase capacity or redundancy, this work list above is done in a factory setting, which not only saves time but also improves predictability of the result. In addition to the quicker initial install, modular UPSs offer the benefit of being able to scale capacity over time with minimal work that takes only hours vs. the days or weeks for the wiring, cabling, and commissioning work to add new UPSs to a non-modular design.

Efficiency impact

The energy efficiency of a UPS is dependent on the load it operates at. And since adding redundancy means adding extra (spare) capacity, redundancy can have an impact on efficiency. Assuming a load of 80% of the 1000 kW rated capacity, which is a typical threshold data center operators set, the UPS configurations analyzed in this paper would operate at 800kW of load. **Table 2** illustrates the implication on load percentage for each configuration, given this assumed load.

Table 2
Impact of configuration on UPS loading

Configuration	Module (unit) capacity	System capacity	Load percentage
1N configuration	250 kW x 4	1000 kW	80%
Internally "modular" redundant	250 kW x 5	1250 kW	64%
Parallel redundant (2+1)	500 kW x 3	1500 kW	53%

The efficiency of any particular UPS at low loads, however, varies from manufacturer to manufacturer, and model to model, and should be investigated as part of the planning process. **Figure 6** illustrates two UPS curves – one that has a much lower efficiency at light loads than full load (left image), and the other that has a fairly flat curve (right image). A UPS such as the one on the left has greater fixed losses, which causes the efficiency to drop at lighter loads, and adding redundancy, in this case, comes at an electricity cost penalty. For a UPS such as the one on the right, adding redundancy will have negligible impact on energy cost. In fact, its optimal efficiency is at a load range of 40-60%. White Paper 108, Making Large UPS Systems More Efficient, provides more background on efficiency curves and the impact operating points have on energy. In addition, a TradeOff Tool (UPS Efficiency Comparison Calculator) is available to help contrast two different UPS curves to see the efficiency and electricity cost implications. **When energy cost is an important decision criterion, it is important to evaluate the UPSs at the expected operating load. The more redundancy added into the configuration, the lighter the operating load percentage.**

Figure 6
Relationship of efficiency and load %. The UPS on the left becomes much less efficient at light loads, the one on the right has a flat curve.

Efficiency curve with high fixed losses

Efficiency curve with low fixed losses

Life Is On   Schneider Electric
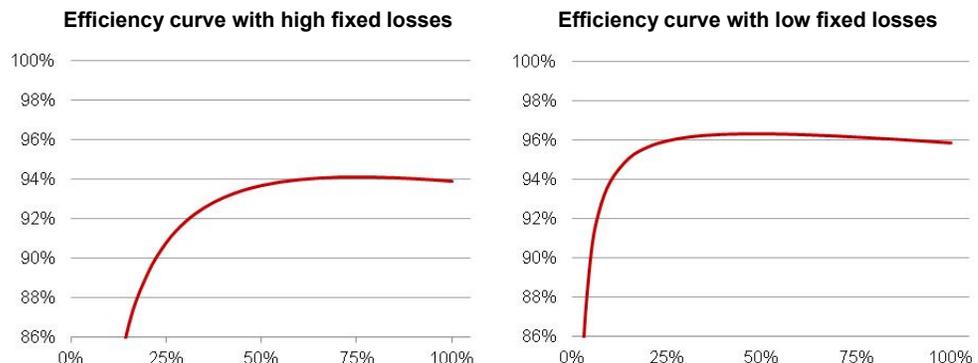
Risk tolerance

Every data center has its own level of risk tolerance based on the criticality of the applications it supports. As we mentioned earlier, more and more fault tolerance is being driven to the IT layer, through technologies like virtualization and hyper-convergence. Based on the IT technologies deployed, an understanding of hard-ware downtime costs to the business (both quantitative and qualitative), the cost premium(s) for the different UPS configurations, and the availability improvements, a decision can be made as to the appropriate level of UPS redundancy.

The cost analysis showed that there was a small premium (6.5%) to go from a 1N design to an internally "modular" N+1 redundant design, and a much larger premium (36.8%) to go from internally "modular" to a parallel redundant N+1 design. In this section, we discuss, qualitatively, the risks of downtime with each of the three configurations. **Table 3** summarizes these risks.

Table 3
Downtime risk comparison

Downtime risk	1N UPS	N+1 internal "modular"	N+1 parallel redundant
Inverter or other power module failure	UPS goes to UPS (static) bypass, no battery backup available	Remaining UPS modules support load, battery backup available	Remaining UPSs support load, battery backup available
DC bus failure	UPS goes to UPS (static) bypass, no battery backup available	UPS goes to maintenance bypass, no battery backup available	Remaining UPSs support load, battery backup available
UPS maintenance	UPS goes to maintenance bypass, no battery backup available	UPS goes to maintenance bypass, no battery backup available	Remaining UPSs support load, battery backup available

INCREASED AVAILABILITY ➤

With a 1N design, any failure within the UPS or its batteries results in a transfer to static bypass. In this mode of operation, a utility failure would bring down the IT hardware.

With internal "modular" redundancy, there is now a spare power module, so that a failure within a single module does not require a transfer to static bypass. Instead, the individual module takes itself offline, while the load remains backed up by the other active modules. The failed module can be replaced later by placing the entire UPS on wrap-around bypass. There are, however, single points of failure in this design. For instance, a failure with the battery system (like a battery breaker tripping) would force a transfer to static bypass, since there is only a single battery bank. Likewise, if the UPS required preventive maintenance, the load would be switched to static bypass or wrap-around bypass, both unprotected from battery.

With a parallel redundant UPS configuration, there is added protection from down-time. Because there are multiple independent UPSs with their own battery strings, the load can remain on protected UPS power during a failure within a single UPS or its batteries. A new risk, however, is introduced, with the controls, communication, and cable impedances to ensure the load is shared across the UPSs. In this paper, we focused on an N+1 configuration where N=2, but depending on the total power needed and the UPS size selected, N may be more than 2. As the N increases, not only does cost and time increase, but reliability may decrease due to increased

challenges in balancing the load current equally through all UPS in all modes of operation.

Human error also has implications on the availability of the various configurations. The more installation field work involved with the design, the greater this downtime risk. See **sidebar**.

Fault tolerance attributes of a UPS

Fault tolerance is what enables a system to continue operating (in this case, supporting the IT load) in the event of the failure of some of its components. With that said, some UPSs are designed with higher levels of fault tolerance than others. When selecting a UPS, it is important to consider the fault tolerance design attributes of the box; especially if the chosen architecture consists of a single UPS frame (as in configuration 1 and 2). Below are examples of fault tolerance design attributes:

- Redundancy of power modules (inverter/rectifier)
- Redundancy of fans
- Redundancy of power supplies in controller
- Redundancy of battery strings
- Redundancy of communication bus
- Redundancy in control system
- Static switch sized for greater than the maximum expected loads to accommodate in-rush / step loads of IT equipment and downstream PDUs

By addressing the critical single points of failure in traditional UPS systems, data centers that once required higher levels of redundancy (like 2N), may be able to rely on these mechanisms to keep critical loads operational. **Figure 7** is an example of a UPS designed with fault tolerance in mind.

<div style="float:left; width:30%;">

Human error

The more field work that occurs (i.e. during installation), the greater the risk of human error leading to downtime; whereas work completed in a factory setting is much more predictable and reliable.

For the three configurations presented in this paper, consider a fault on the output bus (a critical fault that would result in a load drop). With a modular UPS, this bus is internal to the UPS and therefore, factory installed. In the parallel redundant UPS case, the output bus is installed in the field, which increases the downtime risk due to human error.

</div>

Figure 7
Galaxy VX, a UPS family by Schneider Electric with fault tolerance design attributes

There is a general perception that physically separate boxes are required to isolate faults; but it is not always about physical separation, but about the levels of defenses built into the box.

Life Is On | Schneider Electric

Conclusion

As "N+1" becomes a more common UPS architecture for data centers, it becomes more important to understand the tradeoffs of different approaches, so that data center decision makers can make the most educated decision possible given their risk tolerance, budget, and timeline.

In this paper, we discuss the capital costs, deployment timeframe, efficiency, and reliability differences of two common methods of deploying N+1 (and contrast it to a 1N design), to highlight the differences between them. The key conclusions are summarized in the bullets below:

- **Cost:** Internal "modular" N+1 redundant UPS configurations are a capital cost premium of 6.5% over a 1N designs. Parallel redundant N+1 configurations come at a capital cost premium of 36.8% over internal "modular" redundant configurations.

- **Deployment timeframe:** Parallel redundant configurations take about 25-30% longer to deploy than 1N or internal "modular" N+1 configurations. This is the result of additional field work to setup, configure, and ensure communication between separate units. Adding capacity over time also takes longer with parallel redundant configurations.

- **Efficiency**: Redundancy impacts the operating load percentage of a UPS, and for some UPSs, this means an impact on efficiency and electricity cost. Many UPSs today, however, are designed with very flat efficiency curves (lower fixed losses), where the efficiency peaks at part load. This makes this impact negligible.

- **Risk tolerance**: Parallel redundant configurations will provide greater availability to data center loads than internally "modular" configurations. Internal "modular" redundant designs fall between 1N and parallel redundant designs. When selecting a UPS, it is important to consider the design attributes that result in a more fault tolerant UPS.

Internal "modular" redundancy provides significant gains in risk avoidance for a small cost premium and relatively no impact on efficiency and timeline. Parallel redundant UPSs offer even greater risk avoidance but comes at a steeper cost and deployment time premium. In the end, it is the role of the decision maker to weigh these tradeoffs to select the design with the right level for their business needs.

✎ About the author

Wendy Torell is a Senior Research Analyst at Schneider Electric's Data Center Science Center. In this role, she researches best practices in data center design and operation, publishes white papers & articles, and develops TradeOff Tools to help clients optimize the availability, efficiency, and cost of their data center environments. She also consults with clients on availability science approaches and design practices to help them meet their data center performance objectives. She received her bachelor's of Mechanical Engineering degree from Union College in Schenectady, NY and her MBA from University of Rhode Island. Wendy is an ASQ Certified Reliability Engineer.

Life Is On | Schneider Electric

Resources

Comparing UPS System Design Configurations
White Paper 75

Making Large UPS Systems More Efficient
White Paper 108

Electrical Efficiency Modeling for Data Centers
White Paper 113

**Browse all
white papers**
whitepapers.apc.com

UPS Efficiency Comparison Calculator
TradeOff Tool 5

**Browse all
TradeOff Tools™**
tools.apc.com

Contact us

For feedback and comments about the content of this white paper:

Data Center Science Center
dcsc@schneider-electric.com

If you are a customer and have questions specific to your data center project:

Contact your Schneider Electric representative at
www.apc.com/support/contact/index.cfm

Life Is On | Schneider Electric