

Maximizing UPS Availability

A comparative assessment of UPS designs and deployment configurations for the high-availability data center

*By Chris Loeffler
Product Manager, BladeUPS & Data Center Applications
Eaton Corporation*

Executive summary

Uninterruptible power systems (UPSs) play a vital role in ensuring IT reliability. As a result, their own reliability is a crucial consideration too. Any time a UPS fails, mission-critical electrical loads are potentially at risk.

What, then, can organizations do to optimize UPS availability? As this white paper shows, the conventional answers to that question are often not the best ones. UPS reliability is ultimately less a function of UPS design—such as the differences between line-interactive and double-conversion products—than of overall power system design. In the end, the surest way to increase UPS availability is to focus on minimizing repair time and maximizing redundancy, both inside your UPSs and across your power protection scheme as a whole.

In addition, this white paper also explains why, contrary to popular belief, modular UPS designs provide superior availability even though they typically contain more parts that could potentially fail.

Table of contents

The Myth of MTBF	2
UPS Designs and Internal Power Paths	3
Strategies for Increasing the Availability of UPS Power Paths	5
Increasing Availability by Deploying Paralleled UPSs	6
Parallel UPS Architectures	7
How Batteries Impact Reliability	10
Conclusion: Six Key Steps to Maximizing Power System Availability	11
About Eaton	12
About the Author	12

The myth of MTBF

Historically, mean time between failures (MTBF) has been a key metric that UPS manufacturers use to measure and express reliability. In truth, however, MTBF is generally a poor means of predicting UPS availability.

To understand why, consider a UPS with an MTBF of 200,000 hours. A layperson might expect such a device to experience one failure in 200,000 hours—or 23 years—of operation. In reality, however, UPS manufacturers can't and don't test their products for 23 years. Instead, they calculate an initial MTBF based on the projected lifespan of the UPS's components. Then, after they've shipped a statistically meaningful number of units, they replace those preliminary estimates with new ones based on actual performance in the field. Those revised numbers can be misleading though. For example, if 2,500 UPSs perform flawlessly over a five-year study period, the result will be an impressively high MTBF rating. But if those systems contain a component with a six-year lifespan, 90 percent of them could fail in the next year following the study period.

In addition, there is no universal standard for measuring MTBF. For years, most government agencies have required manufacturers to provide calculations based on the latest revision of the MIL-HDBK-217F handbook, while many commercial customers have adopted the Telcordia (Bellcore) SR-332 process. More recently, the technology industry has concluded that these measurements, while helpful, should not be the only way manufacturers grade a product's reliability. As a result, manufacturers today increasingly focus on *Design for Reliability* (DFR) as well. Unlike past standards, which concentrate on individual electronic components and their relationship to the circuits used in the product's design, DFR methodologies pay greater attention to a product's intended and expected use under varying conditions.

Still, at the end of the day there remains no one standard for measuring how a UPS performs its mission, which is keeping connected loads powered. As a result, it's nearly impossible to compare one UPS manufacturer's MTBF figures to another's.

Availability offers a somewhat more realistic measure of critical power backup systems. Given the vital role that UPSs play in the data center, the ability to replace aging or failed parts rapidly is crucial. Availability combines MTBF with a second metric called mean time to repair (MTTR) that measures the time required to acknowledge a problem, respond to it and complete a repair.

$$\text{Availability} = \frac{\text{MTBF}}{\text{MTBF} + \text{MTTR}}$$

Availability is typically expressed as a number of "nines" representing the percentage of time over a year's worth of use that a given system is operational. For example, a UPS with an MTBF of 500,000 hours and a MTTR of four hours would have an availability of .999992, or 99.9992 percent ($500,000 \div 500,004$). That translates to an expected downtime of 4.2 minutes per year.

Still, though it's a better gauge of reliability than MTBF numbers alone, availability is flawed in important respects. In particular, it fails to account for time spent on routine service functions. If a system has to be taken down once per year for inspection, recalibration or general maintenance, its actual operational availability will be lower than the formula above suggests.

UPS designs and internal power paths

Though it can increase costs, building multiple power paths into a UPS ensures that if system components such as the rectifier, inverter or internal backup battery fail, no interruption of power to the critical load will result.

There are four basic UPS designs:

- **Standby UPSs** allow IT equipment (ITE) to run off of normal utility power until the UPS detects a problem, at which point the system provides protection against power outages. However, some standby systems offer partial protection against under- or over-voltage conditions, providing only limited use of battery power. Thus, while standby UPSs are highly efficient and economical, they sometimes provide incomplete protection from power problems.
- **Line-interactive UPSs** typically regulate utility voltage up or down as appropriate before allowing it to pass through to protected equipment. However, they must use the battery to guard against various frequency abnormalities, as well as blackout conditions.
- **Double-conversion UPSs** isolate critical loads from raw utility power completely, ensuring that IT equipment receives only clean, reliable electricity. Double-conversion designs are less energy-efficient than standby and line-interactive designs, so they dissipate greater amounts of heat into the data center or equipment rooms.
- **Double-conversion UPS with multi-mode operation** normally operate in a high-efficiency mode that saves energy and money. When power conditions warrant, however, they switch automatically to the greater protection offered by double-conversion mode. In addition, most multi-mode UPSs utilize a modular building block design that increases availability by reducing the time required to perform maintenance and repairs.

Each of these UPS designs features different internal power paths. A **standby UPS** typically has two power paths that are both served by a single power switch. Thus, if the power switch fails, IT equipment loses power. Most standby systems are limited to less than 2 kVA, so a failure generally impacts only a few pieces of IT equipment.

Typical Stand by

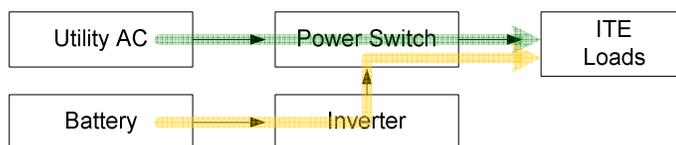


Figure 1. In a typical standby UPS, IT equipment (ITE) loses power if the power switch fails

A **line-interactive UPS** typically has two fully independent power paths, one of which utilizes a power interface. If the power interface fails, the UPS can operate in battery mode long enough to gracefully shut down connected equipment. Some premium line-interactive systems also have a static bypass path that automatically bypasses internal component failures in the UPS and connects IT equipment directly to utility power.

Typical Line Interactive

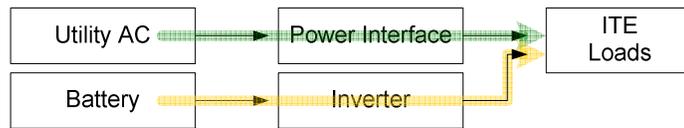


Figure 2. Power paths in a typical line-interactive UPS

Most **double-conversion UPSs** have two power paths—one from utility power or a generator and one from battery power—as well as:

- An *automatic static bypass switch* that bypasses a failed rectifier or inverter and passes utility power directly to IT equipment
- A *manual maintenance bypass device* that permits technicians to service the system without interrupting power to protected loads

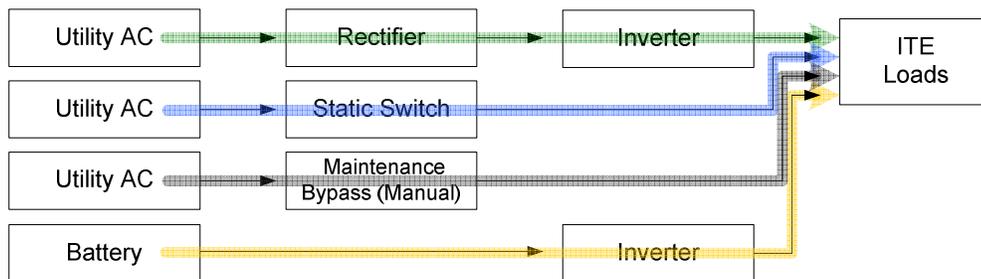


Figure 3. Power paths in a typical double-conversion UPS

Some **double-conversion UPSs with multi-mode operation** have the same power paths as a typical double-conversion UPS plus an *automated* maintenance bypass device that automatically bypasses the inverter whenever UPS electronics are being repaired or serviced. Additionally, if deployed in a modular redundant design, multi-mode UPSs can automatically select whether or not a load needs to be bypassed, ensuring that connected systems stay on backed-up UPS power during almost any service procedure. The result is improved MTTR and reduced risk of downtime or accidental power interruption during maintenance and repairs.

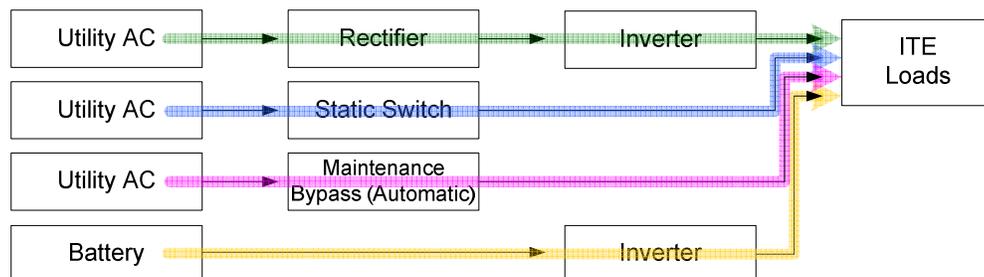


Figure 4. Power paths in a high-efficiency, multi-mode UPS

Strategies for increasing the availability of UPS power paths

There are several ways to increase the reliability of the power paths of a UPS:

- **Add parallel battery strings.** Equipping a UPS with a single string of series-connected batteries can dramatically increase risk of load loss. Say, for example, that a large UPS has 40 batteries connected in a series (+ of the first battery to – of the next). If a problem occurs in any of those batteries, the entire string will probably fail, causing the UPS itself to fail. Adding another 40 batteries and then tying the most positive and most negative points together gives you two parallel strings of batteries. If either string fails, the UPS can typically run for a limited time on the other string until either a backup generator comes on line or load equipment is shut down gracefully.

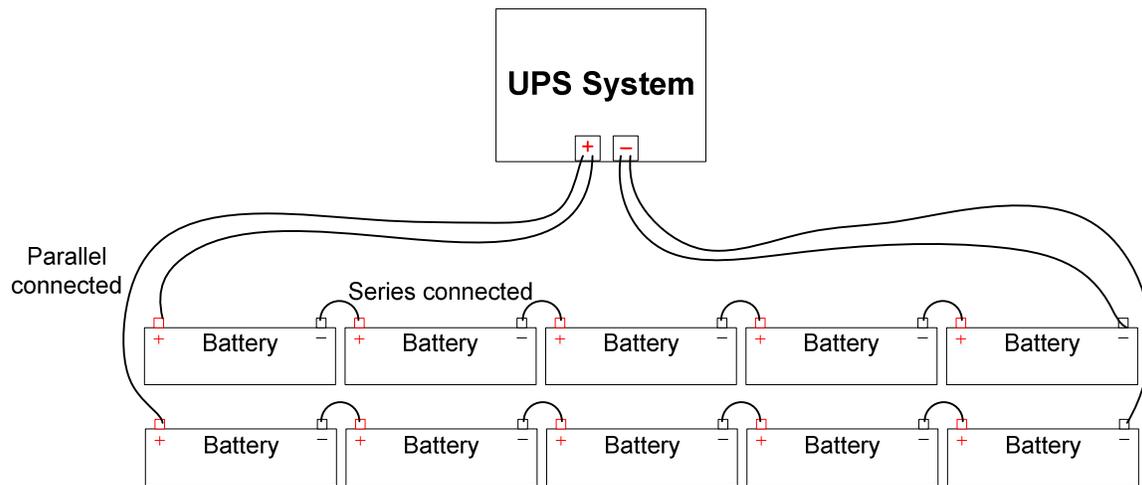


Figure 5. Equipping a UPS with a parallel battery string significantly reduces the chances of a battery failure causing a UPS failure.

- **Install a generator.** Batteries are a temporary source of energy. During a long utility outage, even the best battery strings will fully discharge. A generator provides a second source of backup power during prolonged outages.

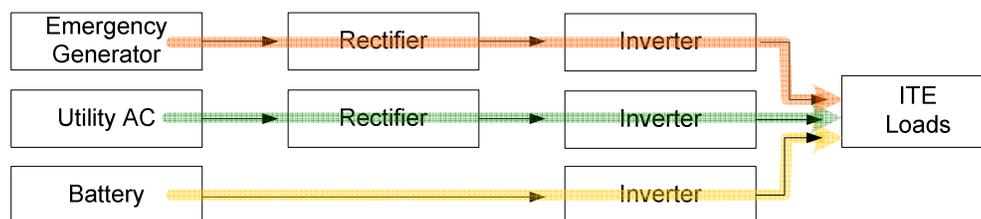


Figure 6. UPS power paths with an emergency generator

- **Ensure that the UPS has an automatic static bypass switch.** This device bypasses the rectifier and inverter, directly connecting utility or generator power to IT equipment in the event of an internal failure in the UPS, or a severe overload or short circuit in the loads supported by the UPS. A static bypass switch typically transfers power without disruption to IT equipment in as little as 3 to 8 milliseconds during a fault condition.

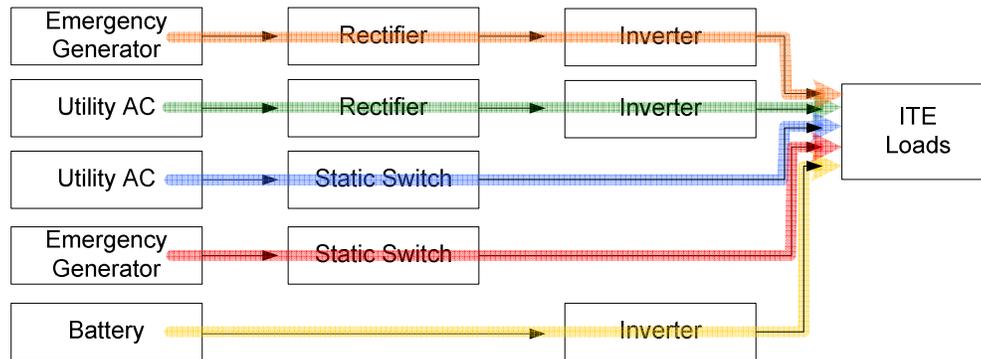


Figure 7. UPS power paths with an internal static switch

Increasing availability by deploying paralleled UPS

The logic of redundancy applies as much to power protection schemes as it does to UPS designs. Building multiple power paths into a power design inherently increases reliability.

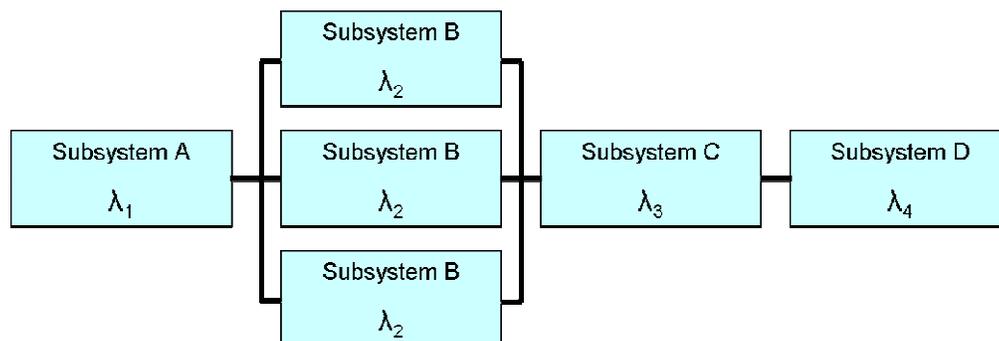


Figure 8. System and subsystem reliability. Source: United States Department of Defense.

Figure 8 illustrates two simple but important points: Power path components that are serially connected, like subsystems A, C and D, weaken overall reliability. Power path components that are parallel redundant, like subsystem B, improve overall availability. That's because a failure in subsystems A, C or D could take down the entire power path. Subsystem B, by contrast, has three parallel components. If one of them goes offline, the others compensate for it, keeping the system as a whole functioning.

Put differently, a power chain is only as strong as its weakest link, so adding more redundancy at each point in the chain increases its overall reliability. As a result, the most reliable power delivery systems typically feature multiple, independent power paths all the way from energy source to load, with as little overlap as possible. In this type of configuration, neither component failures nor routine maintenance procedures result in IT equipment shutting down.

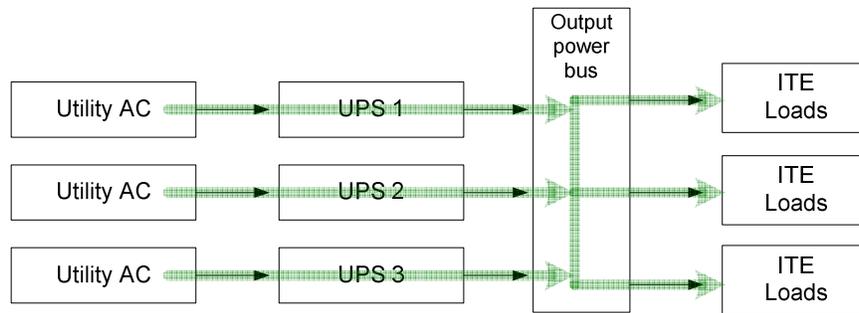


Figure 9. Creating multiple power paths all the way from utility mains to UPSs to IT equipment improves reliability by increasing redundancy.

Parallel UPS architectures

In the UPS industry there are several methods typically utilized in deploying parallel systems. Two of the most popular methods are deploying systems in a cascaded (series) parallel architecture or in a fully redundant parallel architecture.

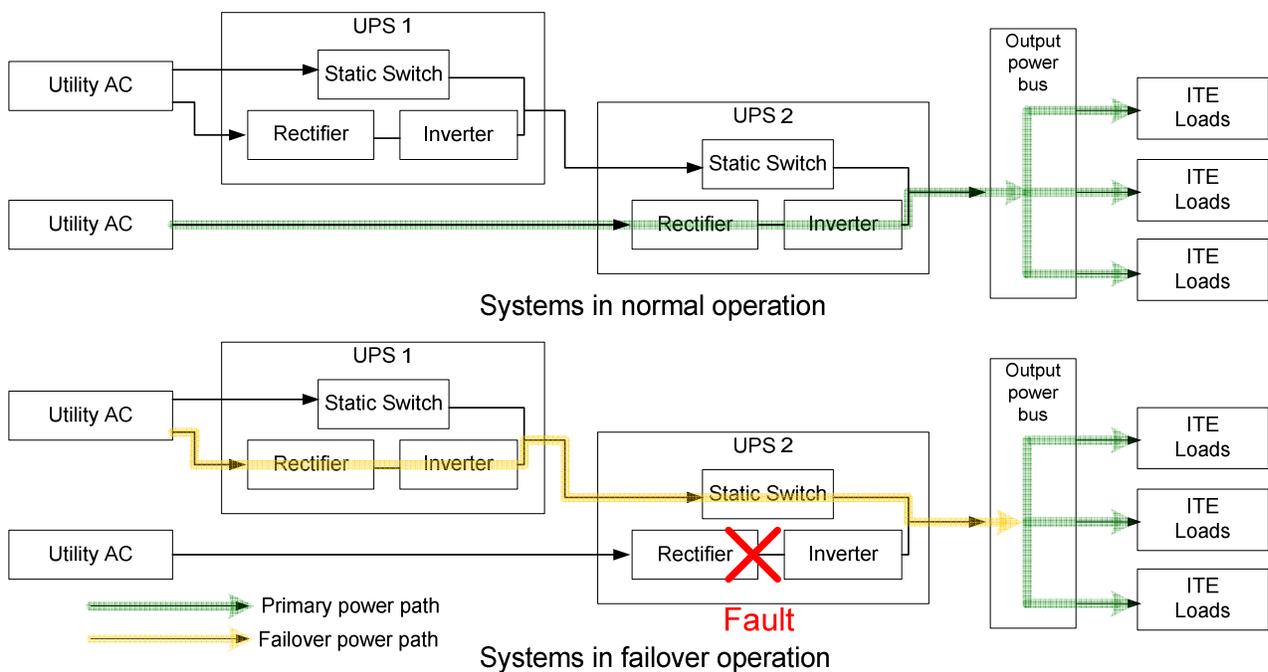


Figure 10. Cascaded parallel systems in standard operation (top) and in a failover condition (bottom).

Cascaded redundant systems are sometimes used in cases where two UPS systems are available to support the base load but they're different models or from different manufacturers, so they can't be paralleled together in a redundant configuration. A cascaded parallel architecture enables you to overcome that limitation.

However, while cascaded parallel systems do offer some limited redundancy, they also require several important events to occur during a failure to ensure that loads stay protected. These events are:

- 1.) Failed system must detect the failure is occurring
- 2.) Failed system must make a safe transfer to its own internal static switch
- 3.) Failed system must remove its failed components from the output power bus
- 4.) Backup system must be able to support full load immediately when requested (load step)

In addition, cascaded parallel systems also impose the expense of operating and maintaining a UPS with no load.

Fully redundant parallel architectures generally offer higher reliability. This depends on how they're implemented, however. Some UPSs said to have a parallel architecture actually have limited parallel components. That is, while they provide limited redundancy if a similar part fails, they lack individual subsystems. If a subsystem fails, then, the entire UPS typically shuts down and needs to be repaired.

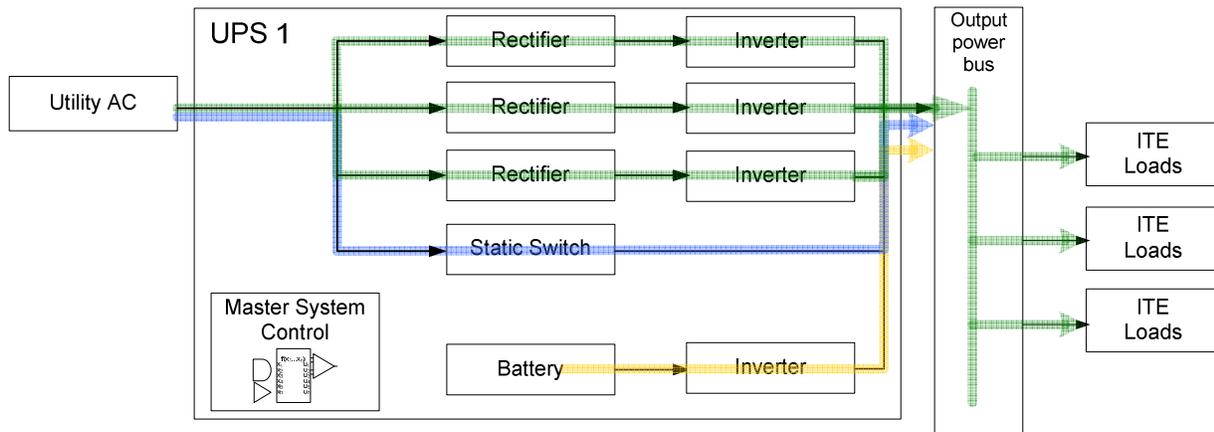


Figure 11. Parallel architecture with some internal redundancy.

Other UPS designs include individual subsystems as well as peer-to-peer paralleling capability, meaning they control themselves rather than utilize a master controller, giving them the highest level of reliability. The idea of any parallel architecture is to eliminate as many single points of failure as possible without adding undue complexity to the design. Thus architectures utilizing individual subsystems and peer-to-peer control offer the most reliable design with the fewest points of failure.

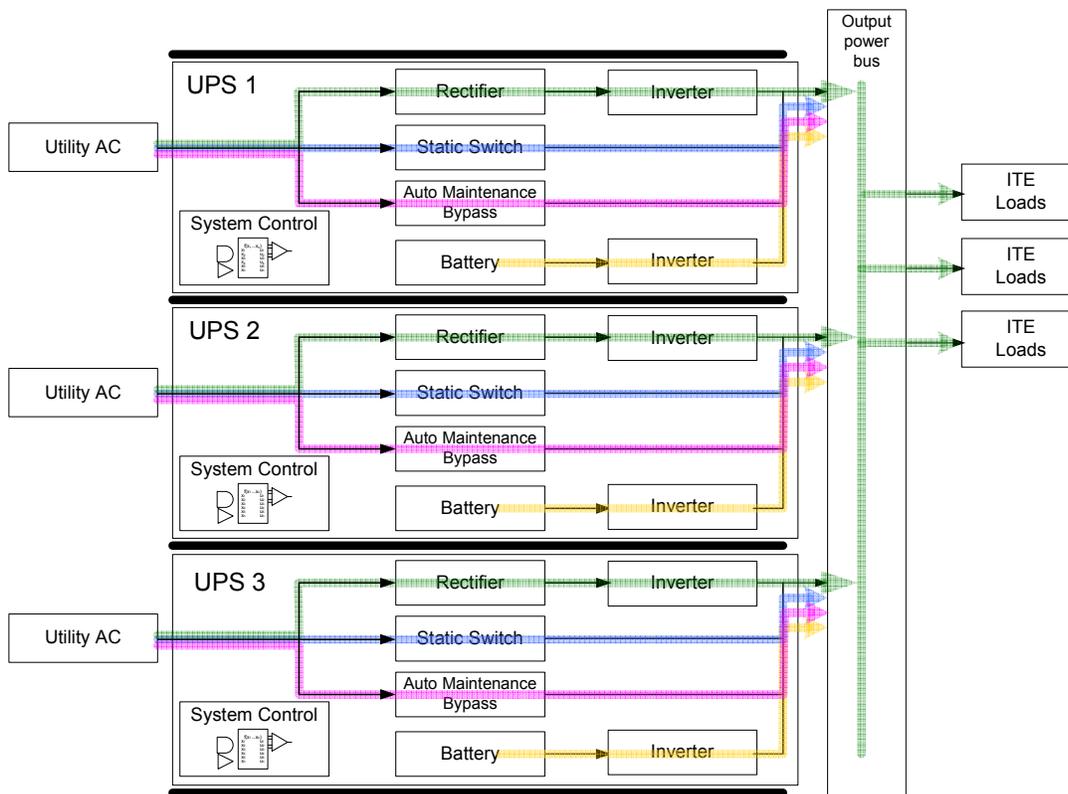


Figure 12. Parallel redundant architecture with peer-to-peer control and individual subsystems for each UPS.

Of course, a parallel redundant UPS configuration with more components and connections also has more potential points of failure and thus a lower MTBF. As a result, IT managers often assume that the fewer UPSs they include in a parallel architecture, the more reliable it will be. However, while adding components to a UPS architecture does eventually reach a point of diminished returns, a conservatively designed system with more UPSs ultimately delivers greater availability than one with fewer UPSs.

To understand why, consider two sample parallel redundant approaches to protecting a 60 kW load. The first features two traditional 60 kW UPSs. The second uses six newer 12 kW UPSs composed of modular building blocks.

Now imagine the impact of a hardware failure on both configurations:

- The 60 kW UPSs can only be repaired by a trained professional. Even if that person arrives within four hours, total downtime for the impacted unit would likely be six to eight hours. Moreover, if the service provider doesn't have easy access to replacement parts, downtime could easily extend past 24 hours. Throughout that period, IT equipment would be at heightened risk due to lack of UPS redundancy.
- The 12 kW UPSs, by contrast, come with hot-swappable electronics and battery modules that end users can replace on their own in minutes, assuming they keep spare parts on hand.

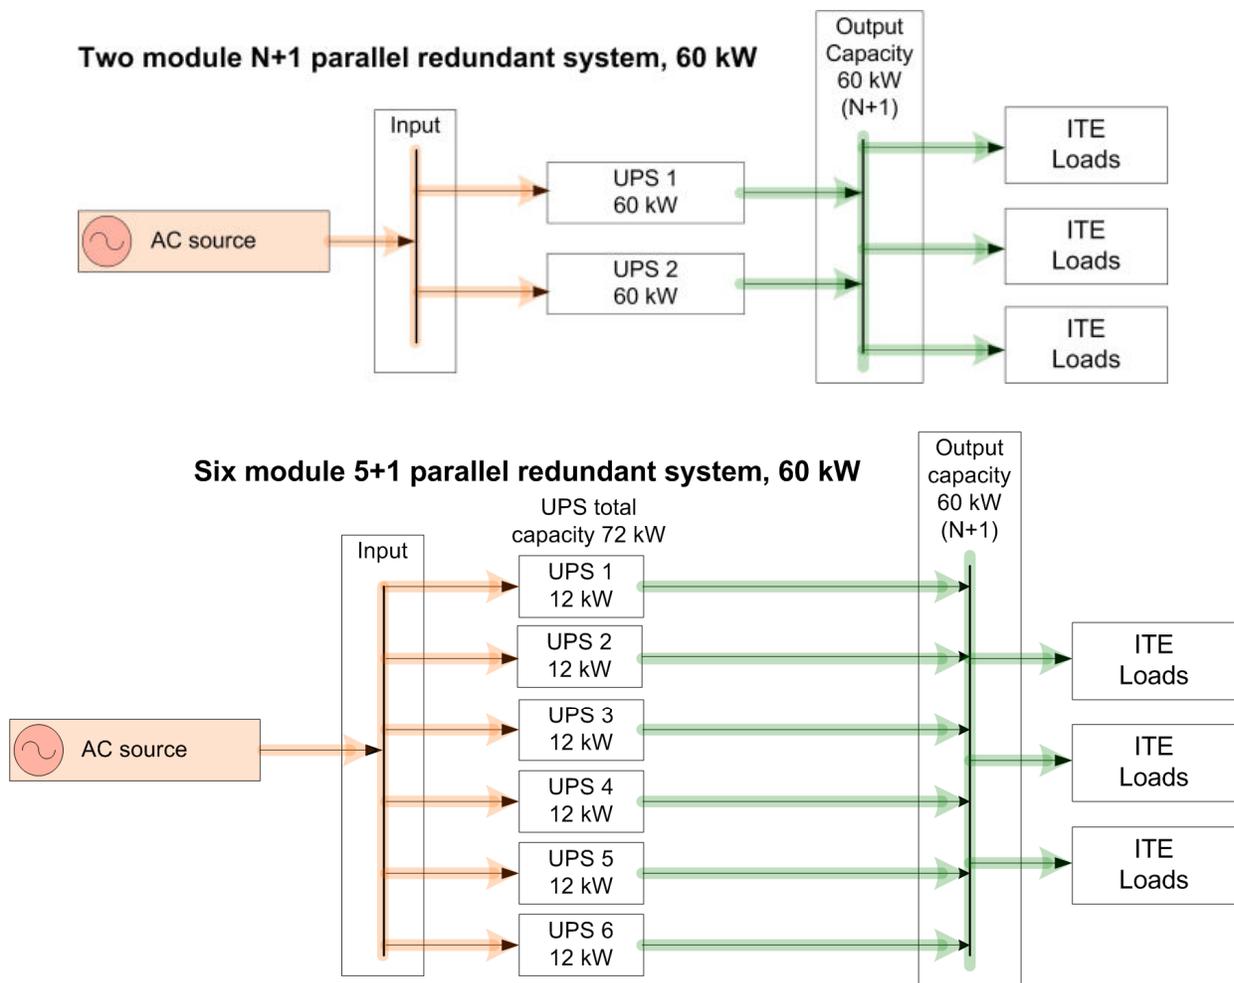


Figure 13. Two options for using parallel redundancy to support a 60 kW load.

Battery-related considerations offer further evidence. The lifespan of a typical UPS battery is four years. Thus, the 60 kW UPS configuration is likely to lose redundancy due to a battery-related issue for at least six hours every four years. The 12 kW UPS configuration, on the other hand, would probably lose redundancy for approximately one hour every four years.

What is true of batteries is true as well of electronic and mechanical components like fans and capacitors, all of which are typically considered wear-out or consumable items in a UPS. Products designed with hot-swappable parts will always experience less downtime. Thus, even though the six-UPS architecture may have a lower part failure MTBF than the two-UPS architecture, its lower MTTR ultimately results in far better availability.

How batteries impact reliability

The design of a UPS dictates how frequently it uses battery power, which in turn affects battery runtime and service life.

Standby UPSs shift frequently to battery mode, which can reduce battery runtime and service life.

Furthermore, the brief interruption of power during those frequent transfers could lock up IT systems, and their typically wide output voltage regulation window may cause the IT power supplies to shut down.

Line-interactive UPSs provide a higher level of protection against power anomalies than standby systems, but must typically resort to battery power when transferring between normal and regulated modes or coping with the instabilities of generator startup.

Double-conversion UPSs are easy on batteries. Within broad tolerances for input voltage, the UPS rectifier/inverter combination can regulate output without resorting to batteries. Additionally, transfer from normal to battery mode is instant, so there's no risk of power interruptions freezing IT systems.

With new **high-efficiency, multi-mode UPSs**, battery usage duration and frequency are similar to a double-conversion UPS and in some instances may even be lower. Furthermore, these UPSs operate at up to 99 percent efficiency under normal use. Higher efficiency translates into longer battery runtime and cooler operating temperatures, which both extend battery service life.

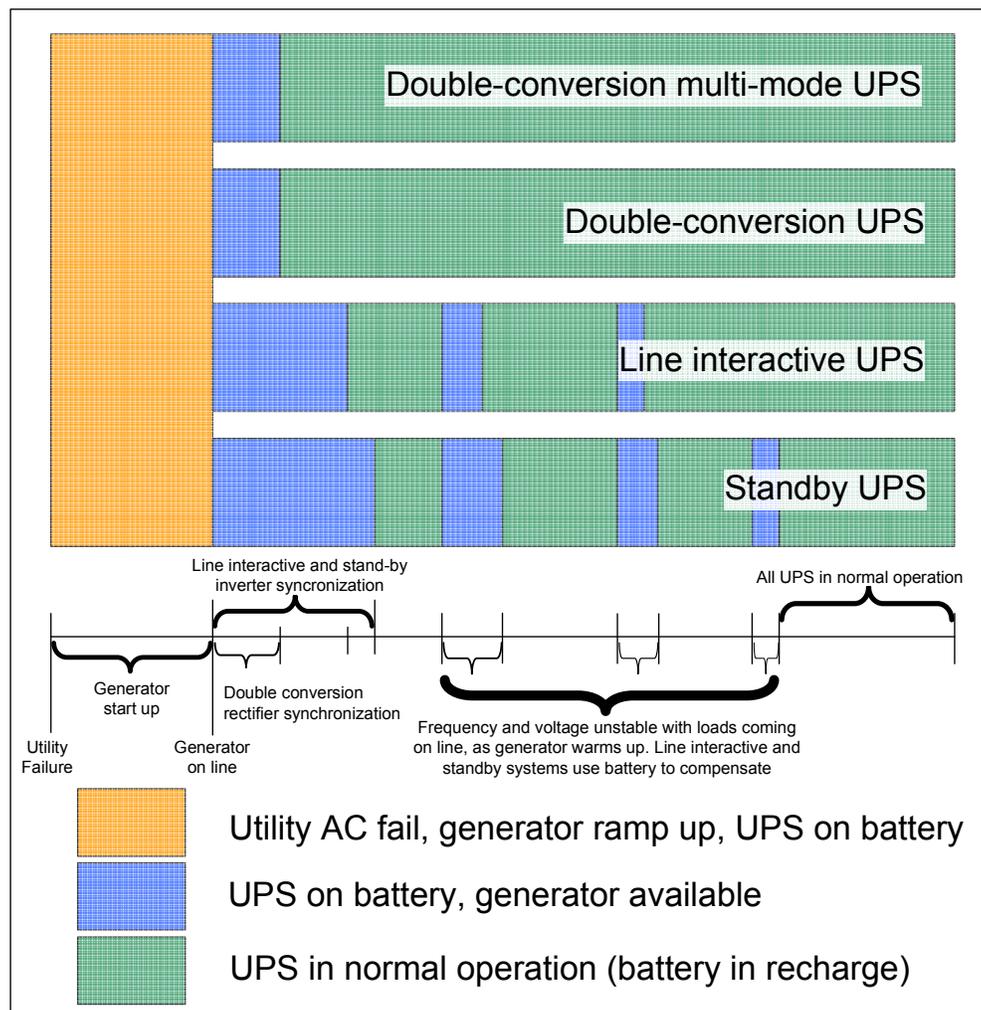


Figure 14. Typical battery usage patterns for various UPS designs.

Conclusion: six key steps to maximizing power system availability

1. **Standardize on a high-quality UPS.** Select a manufacturer with significant experience and a proven in-service success record. Look for designs with internal redundancy of key components, multiple power paths, higher quality components and manufacturing processes that incorporate rigorous quality tests.
2. **Choose UPSs with multiple internal power paths.** Better UPS designs provide multiple power paths for additional redundancy, including such features as a static bypass switch and manual or automated maintenance bypass.

3. **Look for a UPS that is capable to support your IT equipment.** Some low-cost UPS designs may not be able to properly support your load, causing IT equipment to reset, corrupt data or shut down. Double-conversion and high-efficiency, multi-mode UPSs supply conditioned power well within the acceptable voltage and frequency range of the IT and industrial equipment they support.
4. **Deploy redundant, parallel UPSs.** This strategy establishes redundancy of power paths, electronics and battery modules to create the highest reliability protection.
5. **Look for features that improve MTTR.** Select modular system designs and UPSs with easily serviceable parts, such as hot-swappable batteries and electronics. Ultimately, MTTR has a far more profound impact on availability than MTBF.
6. **Select a UPS design that minimizes battery use.** Battery runtime and service life are shorter in UPSs that use battery power frequently. Double-conversion and high-efficiency, multi-mode UPSs make less use of batteries, which typically extends battery life.

About Eaton

Eaton Corporation is a diversified power management company with 2009 sales of \$11.9 billion. Eaton is a global technology leader in electrical components and systems for power quality, distribution and control; hydraulics components, systems and services for industrial and mobile equipment; aerospace fuel, hydraulics and pneumatic systems for commercial and military use; and truck and automotive drivetrain and powertrain systems for performance, fuel economy and safety. Eaton has approximately 70,000 employees and sells products to customers in more than 150 countries. For more information, visit www.eaton.com.

About the author

Chris Loeffler is a product manager for Eaton Corporation, specializing in data center power solutions and services. With more than 19 years of experience in the UPS industry, he has overseen product management of more than 20 UPS products for data center and industrial applications. Mr. Loeffler has held a variety of positions with Eaton, including roles in service engineering, application engineering, program management, and more than 10 years within product management. Mr. Loeffler has authored a number of articles for trade publications and written several white papers on energy efficiency, virtualization and cloud computing in the data center. He has also written articles on various UPS topologies for data center and industrial applications.

Tutorials on demand

Download Eaton white papers to learn more about technology topics or explain them to customers and contacts. Maintenance bypass, paralleling, UPS topologies, energy management and more are demystified in complimentary white papers from our online library: www.eaton.com/pq/whitepapers.