

# Parallel UPS configurations

*Connecting multiple UPS modules for added capacity or redundancy*

*By Ed Spears*

*Product Manager, Eaton Power Quality Solutions Operation*

## Abstract

Increasingly, organizations are finding that the risk of running off straight utility power—even briefly—is too great to ignore. So they deploy multiple UPS modules to ensure conditioned power even if one UPS fails.

In *paralleling*, two or more UPSs are electrically and mechanically connected to form a unified system with one output—either for extra capacity or redundancy. In an N+1 redundant configuration, there is at least one more UPS module than needed to support the load. As a conjoined system, each UPS stands ready to take over the load from another UPS whenever necessary, without disrupting protected loads.

Today's firmware-based paralleling offers particular advantages, compared to traditional paralleling approaches. For one, there is no system-level single point-of-failure. With a peer-to-peer control strategy, each UPS module operates independently and is not reliant on an external master controller or a complex web of inter-module control wiring.

## Contents

<b>The growing trend toward parallel UPS systems .....</b>	<b>2</b>
<b>How do parallel UPS configurations work? .....</b>	<b>2</b>
<b>Four key challenges in parallel UPS systems.....</b>	<b>4</b>
Master control—who's running the show? .....	5
Synchronizing the output of individual UPSs into shared output.....	5
Balancing the load equally among the paralleled UPSs .....	6
Selective tripping—identifying and temporarily isolating a UPS with a problem .....	7
<b>Design for reliability .....</b>	<b>7</b>
<b>Deploying parallel UPS systems .....</b>	<b>9</b>
Customization options for large parallel systems.....	9
<b>Other options for establishing redundant UPS protection.....</b>	<b>9</b>
<b>Closing thoughts .....</b>	<b>10</b>
<b>About Eaton.....</b>	<b>10</b>
<b>About the author .....</b>	<b>10</b>

# Parallel UPS configurations

*Connecting multiple UPS modules for added capacity or redundancy*

*By Ed Spears*

*Product Manager, Eaton Power Quality Solutions Operation*

Electronic systems require conditioned, continuous power—and they get it from uninterruptible power systems (UPSs). But what happens if a UPS is offline for any reason? In that case, the UPS switches to an internal bypass path, and power bypasses the internal power quality circuitry inside the UPS. Critical loads run off utility power until the UPS can be brought online.

Increasingly, organizations are finding that the risk of running off straight utility power—even briefly—is too great to ignore. So they deploy several UPS modules to ensure conditioned power even if one UPS fails.

In *paralleling*, two or more UPSs are electrically and mechanically connected to form a unified system with one output—either for extra capacity or redundancy. In an N+1 redundant configuration, you would have at least one more UPS module than needed to support the load. As a conjoined system, each UPS stands ready to take over the load from another UPS whenever necessary, without disrupting protected loads.

Let's take a closer look at parallel UPS architectures—how they work, what challenges must be overcome in establishing parallel configurations, how modern paralleling technology enhances availability, and what difference it makes in your power protection scheme.

## The growing trend toward parallel UPS systems

Redundant UPS configurations were once relatively rare. Organizations balked at the expense of buying two UPSs to do the work of one. Only the most substantial organizations—or those with the most critical power requirements—made the investment.

That has changed. Data center managers and facilities managers have come to the conclusion that running off raw utility power, even briefly, represents unacceptable risk. The cost of downtime is now so high that even small data centers can justify the cost of redundant UPSs. In fact, redundancy is a requirement of data centers at Tier II and above by Uptime Institute standards.

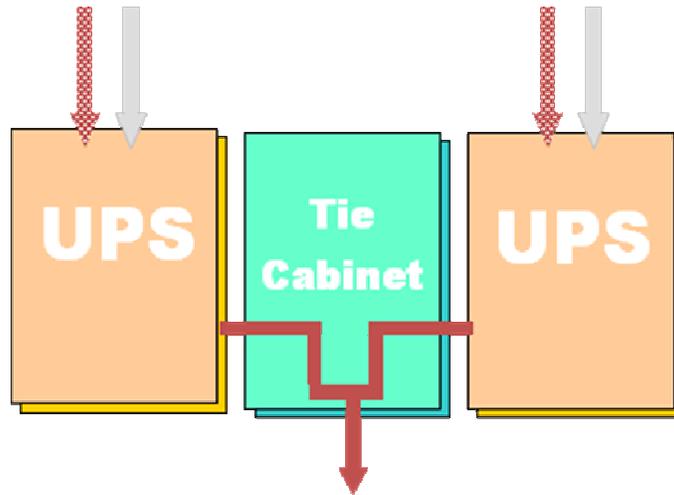
As a result, parallel UPS configurations are becoming commonplace. At least 50 to 60 percent of large UPS systems (300 kVA and up) are configured as parallel systems. Ten years ago, it wasn't even possible to parallel smaller systems (in the neighborhood of 10 kVA), but now up to 40 percent of these smaller systems are paralleled—particularly in Europe and Asia.

## How do parallel UPS configurations work?

On the surface of it, the concept of paralleling UPSs for redundancy is simple enough. Multiple UPS modules are linked to perform in unison, sharing the critical load among them via a common output, and ready to take over for any other module if necessary. In an N+1 configuration (a typical redundancy arrangement), there would be sufficient spare capacity to support the load if any one module became unavailable.

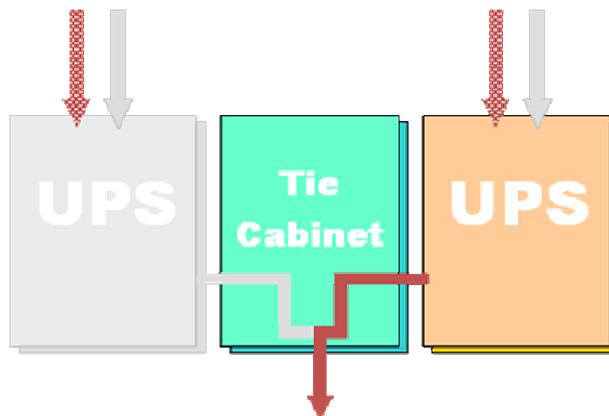
For example, you could protect an 800 kVA load by deploying three 400 kVA UPS modules. During normal operation, the three modules would each carry one-third of the total 800 kVA load. If one module went offline, the remaining two modules would have sufficient capacity to support the load.

The diagram shows a typical parallel configuration with two three-phase UPS modules. In normal operation, AC power flows from the utility source to each UPS—one input into the rectifier and one into an internal bypass. The UPS converts incoming AC power to DC and then back to AC, then sends this clean power to a tie cabinet, where outputs from both UPSs are merged into a single output to protected loads.



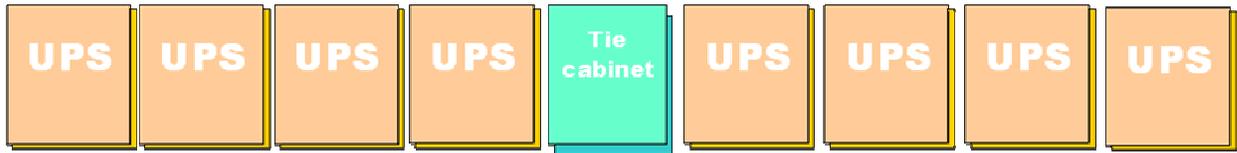
**Figure 1. In normal parallel operation, both UPS modules contribute equally to shared output.**

Should a failure of any kind occur with either module, the critical load is still UPS-protected. Internal diagnostics immediately isolate the faulty UPS module from the critical bus while the other UPS assumes the full load.



**Figure 2. If either UPS module becomes unavailable, the remaining module assumes the load.**

A parallel configuration is not limited to two UPS modules. It frequently includes up to four modules. With some newer UPSs, you can parallel as many as eight modules.

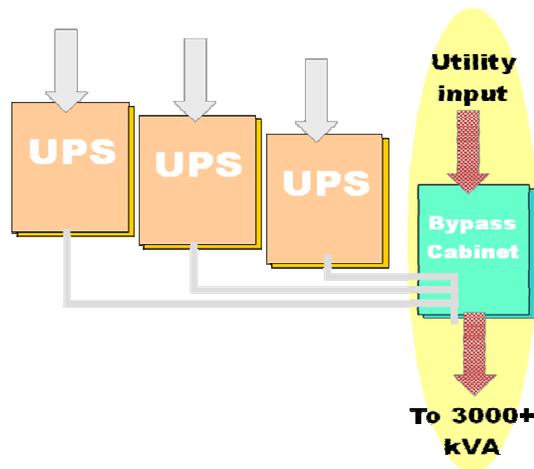


**Figure 3. Up to eight UPS modules can be paralleled into a single system.**

During a utility failure, each UPS module is supported by its battery system and can continue operating for minutes or hours, depending on how much battery runtime has been provisioned. You can (and should) provision separate battery backup for each UPS, for even greater backup protection.

The configuration shown in Figure 4 has a bypass cabinet rather than the standard tie cabinet. When many UPSs are linked in parallel, the load they collectively support could exceed the capacity of the internal static switch and bypass circuit in any one UPS. The bypass cabinet, with its own static switch, provides an alternate route for power during a failure—an automatic and instant *wraparound bypass*.

Such an event would be rare. The wraparound bypass would be activated only if all three UPSs were unable to support the load. Perhaps a short circuit caused an extraordinary overload that exceeded the capacity of all three modules together. The system would identify a failure on the critical bus and transfer to bypass mode with virtually no interruption.



**Figure 4. Power flows to critical loads, even if all three modules were offline.**

### Four key challenges in parallel UPS systems

As soon as you connect multiple AC power sources into a unified, parallel system, there are four key challenges to address:

- Controlling how the separate UPSs should cooperate as a unified system
- Synchronizing the output of each UPS so it can flow into a shared output
- Balancing the load equally among all UPSs in the parallel configuration
- If trouble occurs, identifying and temporarily decommissioning the UPS with the problem

These issues can be complex, and they must be managed in a way that does not compromise the high reliability for which UPSs are paralleled in the first place.

Parallel UPS configurations have been in use since the 1970s for huge defense and military installations, the U.S. Federal Aviation Administration and many other large industrial, commercial, government and healthcare facilities. After 40 years of in-service experience, paralleling technologies have evolved quite a bit in the way they manage all four challenges: master control, synchronization, load balancing and “selective tripping.”

### **Master control—who’s running the show?**

#### **In a conventional parallel UPS system...**

Control is usually managed with a master-slave arrangement that relies on a veritable spiderweb of control wiring between UPS modules. A master controller serves as the brain of the entire system, determining how UPS modules synchronize their outputs, how they share loads, and where they get this information.

Good parallel systems will have redundant master controls, and the better ones will have a “moving master, sliding slave” configuration—a round-robin of rotating leadership. If the master controller fails, one of the other UPS modules is designated as the new master, and the other UPS modules must look to the new source for their commands. If multiple modules fail, a new master is designated based on some pre-negotiated scheme.

This process of shifting leadership is fraught with peril, and it is easy to see why. The change of command must be smoothly executed by every module. A simple glitch with inter-module wiring can create havoc. For example, if control wiring to UPS #5 goes down, paralleling won’t work for UPSs #1 through #4 either.

When customers look closely at parallel UPS configurations, they immediately identify this process as the Achilles heel. The more closely they look at it, the more they’re appalled at the vulnerability of it. Failure of the control system to automatically switch to a redundant path is the leading cause of failures in mission critical power systems, accounting for more than 30 percent of incidents.

#### **In a more resilient parallel UPS configuration...**

Instead of using a master-slave arrangement, a peer-to-peer control system manages the multiple UPS modules, very similar to peer-to-peer computing networks. Each UPS assesses its own operating parameters and determines how to interface with the others. There is no need for control wiring among the UPS modules. No matter what happens to other modules, the parallel system still functions, because each module contains the intelligence it needs to be a functioning member of the group.

### **Synchronizing the output of individual UPSs into shared output**

UPSs in a parallel configuration must deliver output at a specified voltage and frequency—that is, have their sine waves completely aligned with each other. If their output is not synchronized, the voltage disparity would cause a large and potentially damaging surge current from UPSs to the load.

#### **In a conventional parallel UPS configuration...**

Synchronization depends on a master controller. All UPS modules in the parallel configuration look to that master controller to get their synchronization information. If that master controller goes offline—or there is a glitch in the tangle of inter-module control wires—the whole system is in trouble. UPS modules would get out of sync with each other, causing overload conditions on one or more modules and perhaps triggering those modules to go offline in self-protection.

**In a more resilient parallel UPS configuration...**

Each UPS synchronizes to the bypass source, which is common to all the modules. If one module loses bypass feed, it can still synchronize to another module's bypass. In the absence of a bypass source, the module simply looks at its own output and makes adjustments based on what it finds. The process is completely autonomous and intrinsic, thanks to inventive software and very fast microprocessors. In this design, there is no need to distribute synchronization signals, so the system is not at the mercy of a single potential point of failure.

**Balancing the load equally among the paralleled UPSs**

UPS modules in a parallel configuration should share the load evenly with each other. In a two-module configuration, the load would be distributed 50-50. In a four-module configuration, 25-25-25-25, and so on. Load-sharing ensures that no UPS module is overloaded, nor is any module unnecessarily stressed by suddenly bursting from low load to high load as conditions change.

**In a conventional parallel UPS configuration...**

Load-sharing typically depends on a load-share loop, whereby UPS modules continually communicate their status to each other through a web of control wiring. This communication is complex, critical to the operation of the system—and fragile. If any part of the communication web fails, so does the system.

A common problem in this arrangement is electrical noise. Any time you loop communication wires among UPSs—a high-power, high-frequency environment—they are prone to picking up electrical noise every step of the way. A communication disturbance can be disastrous to the operation of the system, but it is also very difficult to troubleshoot.

What happens when the load becomes unbalanced among UPS modules? Suppose you have a three-module parallel system, and due to a faulty load-sharing mechanism, modules #1 and #2 are each carrying 40 percent of the load, and module #3 is carrying only 20 percent. Chances are, UPS #2 will gradually assume more and more of the load, become overloaded and go offline. The load is then shifted unevenly to the other two UPS modules, which will likely become overloaded, and the whole system switches to bypass—thereby undermining the benefits of a redundant parallel system.

**In a more resilient parallel UPS configuration...**

Load-share control algorithms in internal UPS software maintain load balance by constantly making adjustments in response to variations in output power requirements. Each module conforms to demand and is not in conflict with the others for the load. As with synchronization, this process requires no inter-module communication. The result is true wireless paralleling.

Modern, firmware-based paralleling can maintain load balance even under adverse conditions, such as:

- When modules are synchronized to an alternate source
- When power backfeed results from the removal of a load
- When some but not all modules lose their synchronizing reference
- When the frequency of the alternate synchronization source oscillates

All of these factors will affect the operation of the load-share function and will frequently conflict with one another. The load-share algorithm makes careful selection of priority and gain to take the most beneficial action. Precise load-share control—considering both active and reactive power—is made possible by a digital signal processing technique known as direct digital synthesis to control inverter frequency.

## Selective tripping—identifying and temporarily isolating a UPS with a problem

When UPS modules are paralleled, it can be difficult to identify the root cause of a failure. You might see a drop in voltage on the shared output bus, but which UPS is the culprit? You need to find out which module is causing the problem, and isolate it quickly before it drags down other modules and causes the whole system to switch to bypass mode.

### In a conventional parallel UPS configuration...

A faulty module may be signified by the whole system going to bypass. Though most module failures are benign, a failed inverter IGBT or shorted capacitor may appear as a fault on the critical bus. For this type of failure, it could be difficult to quickly identify the root cause and remove the failed module from the critical bus. The system could go to bypass, leaving loads exposed to straight utility power for as long it takes for service to arrive.

### In a more resilient parallel UPS configuration...

Each module need only look at itself to see if it has failed. An algorithm assesses the difference in current/voltage in each phase and detects failures based on a running record of this information—continually comparing present waveforms with previously recorded waveforms. Based on this high-speed calculation, the unit detects a fault even before typical hardware sensors would detect it. The affected UPS module turns off its inverter IGBT transistors within microseconds (millionths of a second). The result is a *selective trip* that instantly isolates the faulty unit from the system until the problem can be resolved.

Since this process does not require communication links among modules, the module is swiftly removed from the critical bus before the problem can affect critical loads and before the system sees the need to go to bypass.

## Design for reliability

Redundant UPS configurations are a necessity to meet the uptime requirements of a 7/24 world—and paralleling is a key way to maximize that uptime. However, when two or more AC power sources are joined in parallel, you may solve one problem while potentially creating others, for all the reasons described earlier.

Most parallel technologies on the market can adequately meet the needs for synchronization, load-sharing and selective tripping—but if you look closely at how they perform these functions, you'll see big differences in potential reliability on several key dimensions:

- **Autonomy**—Can the system successfully operate without external controls and monitoring?
- **Complexity**—How many components, connections and negotiated interfaces are required?
- **System wiring**—Is there a complex and vulnerable mesh of communication wires between modules?

This last element has proven to be the weakest link in conventional parallel systems, which require a great deal of control wiring between modules and sometimes between modules and the bypass cabinet. A typical arrangement has a set of wires for synchronizing each phase and neutral, another set of wires for load-sharing, and yet another set of wires to control the selective tripping process. Very quickly this architecture yields a huge bundle of control wires that are devilish to troubleshoot.

For example, consider that load-sharing control wires are typically arranged in a loop from UPS #1 to UPS #2 to UPS #3 to a bypass cabinet and back to UPS #1. Suppose you had a loose connection between UPS #1 and UPS #2. When you connect test equipment to the wires, you would see the problem no matter where you looked, because the communications network is a continuous loop. With hundreds of wires—any one of which could be improperly wired, loose or disconnected—it can take hours or days of painstaking manual effort to identify the source of trouble.

In contrast, a more resilient paralleling design provides the four critical attributes—control, synchronization, load sharing and selective tripping—without a master controller and without control wiring between modules. The result is true wireless paralleling, completely autonomous and intrinsic—without many of the potential points of failure typical in other designs.

Parallel architecture	Autonomy	Circuit complexity	Inter-module wiring	Integration
Master control	Many failure modes	Negotiation needed for multi-master	Critical	Extrinsic
Master synchronization	With passive loop buffer	Master clock	Critical	Intrinsic
Load-share loop	With passive loop buffer	Vector sum	Critical	Intrinsic
State-of-the-art paralleling	ABSOLUTE	NONE	NONE for control	Intrinsic

### Parallel systems for added capacity

Most organizations plan to grow, but when and how much? How much power will you consume next year, or in five years? You don't want to overbuild the power system today for future demands that may or may not materialize. Even if you could justify the cost, the power infrastructure would operate far below capacity and be very inefficient as a result. And you certainly don't want to rip out and replace today's UPS just because next year's moves, adds and changes suddenly double the need for power.

Paralleling provides an excellent solution for matching growth while extending the value of existing UPSs. The architecture to parallel for capacity looks very similar to paralleling for redundancy. Hardware components are the same; there are just small differences in operation.

A system paralleled for capacity allows you to add load until it reaches capacity, then notifies you to add another module. In contrast, a redundant parallel system constantly ensures that there are enough modules to take over the total load if one drops off (N+1). For example, if the parallel system has five 100 kVA modules, the system would issue an alarm if the load exceeded 400 kVA—the load that four of those five modules could support.

## Deploying parallel UPS systems

In a well-designed paralleling design, all you need is two or more compatible UPS modules and an electromechanical tie cabinet that connects the output of those UPS modules together. Ideally, no special circuitry or software is required in the UPSs themselves. That means existing UPSs in the field can become part of a parallel system without retrofitting or replacement.

UPSs of different ratings can have widely different options in terms of number of modules that can be paralleled and maximum overall rating/redundancy. The examples below are typical:

Rating per module	Paralleling opportunity
4 kVA and 5 kVA	Up to 8 kVA or 10 kVA with kit Up to 20 kVA with frame
8–15 kVA	Up to 45 kVA (N+1)
12 kW	Up to 60 KW in a single rack
200–500 kVA or 500–750 kVA	Up to 3300 kVA
10–30 kVA	Up to 90 kVA (N+1)
40–160 kVA	Up to 800 kVA
225–1100 kVA	Up to 3300 kVA (N+1)

With the newest, rackmounted UPSs designed for high-density server environments, no tie cabinet is required. Paralleling is accomplished using a plug-and-play bus structure that mounts easily in the back of the equipment rack.

## Customization options for large parallel systems

In practice, large customers need one-of-a-kind, specialized configurations that match their unique needs for availability and manageability. Many options are available for parallel UPSs, such as:

- Wraparound maintenance bypass, to allow loads to keep running (off straight utility power) even if the parallel system is unavailable, such as during a natural disaster
- Redundant breakers in the tie cabinet, to permit maintenance of the primary breakers without turning the system off
- Separate load bank breakers in the switchgear, to enable use of a load bank to test the UPS system under load while it is isolated from protected loads
- Communication cards and a monitoring system for remote monitoring

## Other options for establishing redundant UPS protection

Redundancy doesn't always require paralleling. There are other options for deploying multiple UPS modules—separate rather than paralleled—to provide an added layer of assurance in the power protection architecture.

For example, separate UPSs can be set up to provide serial redundancy, where even if the primary UPS is offline, its bypass path is protected by another UPS. Or a data center could be divided into separate zones served by separate UPSs, thereby minimizing the impact of a UPS failure. Or separate UPSs could serve either side of dual-corded loads—or source power from different utility substations. Furthermore, any of these options can be set up for duplicate redundancy. However, each option presents some compromises, compared with peer-to-peer configurations described earlier.

## **Closing thoughts**

Nonstop availability of critical systems depends on nonstop performance from the power delivery architecture under all sorts of conditions. Reliable UPS technology is a solid front line of defense, but maximum reliability comes with redundancy.

Where maximum availability and power protection are essential, parallel UPS configurations offer significant advantages over separate, redundant UPSs. Parallel UPS modules can seamlessly share the load and automatically take over for a failed module without disrupting power quality to the critical load, and without unduly stressing the UPSs, other power components or the IT equipment.

The newest paralleling technologies offer particular advantages, compared to traditional approaches. For one, there is no system-level single point-of-failure. With a peer-to-peer control strategy, each UPS module operates independently and is not reliant on an external master controller or a complex web of inter-module control wiring. In fact, no added circuitry or components are required for a UPS module to be switched in to operate in parallel.

## **About Eaton**

Eaton Corporation is a diversified power management company with 2007 sales of \$13 billion. Eaton is a global technology leader in electrical 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 80,000 employees and sells products to customers in more than 150 countries. For more information, visit [www.eaton.com](http://www.eaton.com).

## **About the author**

Ed Spears is a product manager in Eaton's Power Quality Solutions organization in Raleigh, NC. A 29-year veteran in the power systems industry, Spears has experience in UPS systems testing, sales, applications engineering and training—as well as working in power quality engineering and marketing for telecommunications, data centers, cable television and broadband public networks.

You can reach him at [EdSpears@Eaton.com](mailto:EdSpears@Eaton.com).